

Disentangling Human Dynamics for Pedestrian Locomotion Forecasting with Noisy Supervision

Karttikeya Mangalam^{1,3*}, Ehsan Adeli¹, Kuan-Hui Lee², Adrien Gaidon², Juan Carlos Niebles¹

¹Stanford University ²Toyota Research Institute ³University of California, Berkeley

mangalam@cs.berkeley.edu {eadeli, jniebles}@cs.stanford.edu

{kuan.lee, adrien.gaidon}@tri.global

Abstract

We tackle the problem of Human Locomotion Forecasting, a task for jointly predicting the spatial positions of several keypoints on human body in the near future under an egocentric setting. In contrast to the previous work that aims to solve either the task of pose prediction or trajectory forecasting in isolation, we propose a framework to unify these two problems and address the practically useful task of pedestrian locomotion prediction in the wild. Among the major challenges in solving this task is the scarcity of annotated egocentric video datasets with dense annotations for pose, depth, or egomotion. To surmount this difficulty, we use state-of-the-art models to generate (noisy) annotations and propose robust forecasting models that can learn from this noisy supervision. We present a method to disentangle the overall pedestrian motion into easier to learn subparts by utilizing a pose completion and a decomposition module. The completion module fills in the missing key-point annotations and the decomposition module breaks the cleaned locomotion down to global (trajectory) and local (pose keypoint movements). Further, with Quasi RNN as our backbone, we propose a novel hierarchical trajectory forecasting network that utilizes low-level vision domain specific signals like egomotion and depth to predict the global trajectory. Our method leads to state-of-the-art results for the prediction of human locomotion in the egocentric view.

1. Introduction

Pedestrians are one of the most vulnerable and prevalent entities in self-driving scenarios. The ability to predict their dynamics in the near future can assist in making proper decisions for immediate next action the vehicle needs to take. Forecasting human locomotion is useful in several downstream tasks for self-driving cars such as reasoning about pedestrian intent, path planning and reactive control. It is

*This work was completed while the author was at Stanford University.

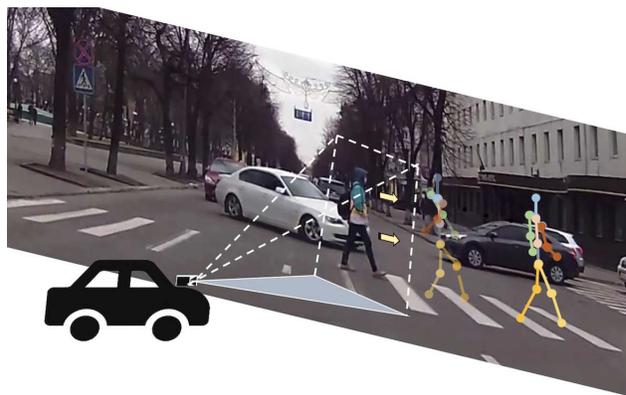


Figure 1: Egocentric pedestrian locomotion forecasting. Locomotion is defined as the overall motion of keypoints on the pedestrian in contrast to predicting just the position (trajectory prediction) or the pose (pose forecasting).

also relevant for social robots that require anticipating future human movements for collision avoidance and navigation. Human dynamics or locomotion can be defined in terms of the joint spatial movement of several keypoints on the human body. Forecasting future human dynamics can lead to foretelling certain undesirable activities like falling, which can then be planned for. It is the final product of a complex interaction between large scale trajectorial motion and finer body limb movements. In previous works, often one of these two critical components is studied. In this work, we propose to predict human locomotion by disentangling the global and local components. An additional challenge to this task is the scarcity of human annotated pedestrian pose datasets in egocentric view. To this end, we use off-the-shelf estimation models to generate ‘noisy’ ground-truth to train our model.

Developing computational methods for modeling human dynamics and forecasting how the pose might change in the future is itself an extremely challenging task. The first level of complexity comes from the inherent multimodal nature of pedestrian locomotion distribution. The space of possi-

ble future locomotion is both complex and uncertain even conditioned on the observed history. Furthermore, in real scenarios, the pedestrians often occlude with each other or other objects in the scene. Moreover, obtaining full annotations of the dynamics (trajectory and pose) is a very intensive task. Therefore, in contrast to previous works [1, 2, 3, 4, 5], applying fully-supervised methods is hardly possible for real-world in-the-wild applications.

In this work, we address the challenge of predicting human locomotion in the wild without explicit human labelled ground-truth supervision. We start with noisy machine generated supervision for forecasting using off the shelf pose detectors like OpenPose [6, 7, 8]. We then propose a forecasting pipeline (Fig. 1) to complete the missing detections and somewhat denoise the detected poses. Furthermore, we identify the complexities associated with learning to forecast human locomotion in the wild in an end-to-end fashion and introduce a disentangling method to overcome this problem. While almost all of the previous work for prediction of human dynamics [1, 2, 9, 10] target the problem of predicting global (*i.e.*, trajectories [11, 12, 13]) and local movements (*i.e.*, human pose and joints motion [1, 3]) separately, we propose to unify the previous works on future pedestrian pose and trajectory prediction under the single problem of locomotion forecasting. To achieve this goal, we propose a sequence-to-sequence pose forecasting model (see Fig. 2) that build upon this decomposition. Lastly, we show that with these advancements our proposed method achieves better results than several other previous works and confirm our hypothesis of disentangling human dynamics for pedestrian locomotion forecasting in the wild.

In summary, our contributions are two-fold. **First**, we propose the pose completion and decomposition modules that uses *disentanglement of global and local motion* to reduce the overall forecasting complexity and enables learning under noisy supervision. **Second**, we present the pose prediction module consisting of a novel egocentric trajectory prediction network based on the Encoder-recurrent-Decoder architecture that utilizes domain specific settings of egomotion to forecast these different granularity streams. They are merged for the final locomotion forecasting.

2. Related Work

Trajectory Prediction: The literature for trajectory prediction can be split into two major categories of work: (1) works that aim to predict human trajectories from the points of view of either other humans in the scene (human-human) [4, 14, 15, 16] or from a top down view [4, 15, 16, 17, 18]; and (2) works that predict the paths of other vehicles on the road from the viewpoint of a car (car-car) [5, 19].

The previous work on human-human trajectory prediction often solves the problem from a top-down prospective *i.e.*, *bird's eye view*. For instance, Social LSTM [4] jointly

predicted the paths of all the people in a scene by taking into account the common sense rules and social conventions that humans typically utilize as they navigate in shared environments. Social GAN [15] defined a recurrent sequence-to-sequence model using Generative Adversarial Networks [20] (GAN) that observes motion histories to predict future behavior with spatial pooling thereby tackling the multimodality associated with motion prediction. In [17], Zou et al. proposed to infer the latent factors of human decision-making process in an unsupervised manner by extending the Generative Adversarial Imitation Learning framework to anticipate future paths of pedestrians (human-car) from the top-down view. In [16], SoPhie introduces an attentive GAN that predicts individuals' paths by taking into account physical constraints, *i.e.*, scene context and other peoples' trajectories. Several work model the human trajectories with social forces [21] and attraction repulsion models using functional objects [22]. Also, approaches using probabilistic modelling of trajectories with Gaussian processes [23] and mixture models [24] have been employed. Among the few works for trajectory prediction in the human's *egocentric view*, [14] predicted trajectories of other pedestrians in an egocentric setting with a tri-stream CNN architecture.

Car trajectory prediction from *egocentric view* were done by explicitly modeling other agents (cars) along with motion with respect to the camera. Yao *et al.* [25] proposed a multi-stream recurrent neural network (RNN) encoder-decoder model for predicting vehicle location and scale using pixel-level observations. DESIRE [19] was proposed to generate several possible trajectories for vehicles using both scene and vehicles, which were then ranked in an Inversion of Control (IOC) framework to choose the most likely one. Similarly, Huang *et al.* [5] introduced an uncertainty-aware method for prediction of self-trajectories.

One of the most crucial tasks for self-driving car or social robots is to predict future dynamics of pedestrians (*i.e.*, car-human task). However, recent work focuses on human-human or car-car prediction, both of which domains lack either egocentric cues like monocular depth cue and large camera motion due to ego-vehicle or pedestrian centered features like human pose and intention.

Pose Forecasting: A number of recent work predicted 3D poses using fully annotated data with exact 3D positions. These works included Structural RNN [3], convolutional sequence-to-sequence model with encoder-decoder architecture [26], and human motion prediction with adversarial geometry-aware based on geodesic loss [27]. Some other works introduced models that can predict poses in both 2D and 3D. Triangular-Prisms RNN (TP-RNN) [1] proposed an action-agnostic pose forecaster, and Matrinez *et al.* [2] proposed a method based on residuals and a careful design of RNNs for human dynamics. These works concluded that zero-velocity is a very strong baseline for pose forecasting.

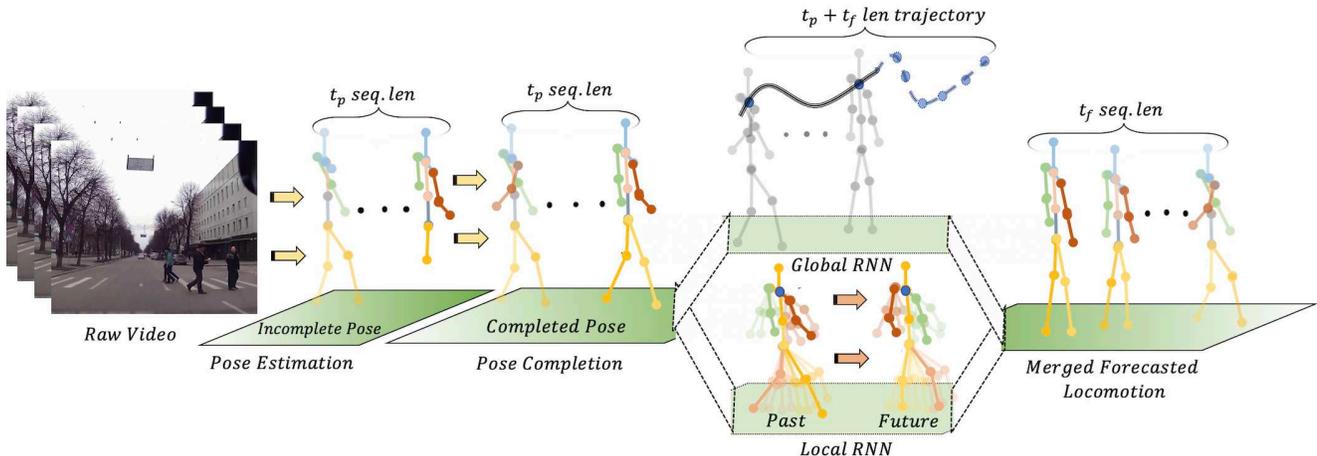


Figure 2: An illustration of the proposed method for human locomotion forecasting with noisy supervision. The “Raw Pose” plane represents the noisy input pose sequence with missing joint detection (Section 3.1). The “Complete Pose” plane denotes the output from the pose completion module with filled in joint positions (Section 3.2). The completed pose is then split into the global and local streams (Section 3.3) which separate concurrent motions. The prediction modules forecast the future streams as outlined in Section 3.5. Finally, these streams are merged to predict future pedestrian locomotion.

Ghosh *et al.* [28] introduced LSTM-3LR for long term forecasting and Zhai *et al.* [29] trained a model to forecast videos of human activities with multiple granularities which they use to forecast pose for future frame generation.

Denosing or Completing Poses: Chang *et al.* [30] propose a Denoising and Reconstruction Networks (DR-Net) for 3D human pose estimation from monocular RGB videos. Ballard *et al.* [31] leverage autoencoders for structured prediction of the human pose. Tekin *et al.* [32] use over complete autoencoders to account for dependencies between joints and use a latent space for 3D pose prediction from monocular images. Trumble *et al.* [33] introduce a convolutional symmetric autoencoder with reconstruction loss to encode skeletal joint positions, which simultaneously learns a deep representation of volumetric body shape.

In contrast, our method for locomotion forecasting is different in three main aspects. *First*, we focus on the egocentric vehicle’s view utilizing domain specific cues whereas the previous works are either on predicting other vehicle’s movements [5, 19] or for forecasting human trajectories from the top-down view [4, 16, 18, 15]. *Second*, all previous works focus on either predicting the pose [1, 2, 3] in the near future or on trajectory prediction [4, 14] whereas we propose a method to unify both for complete human locomotion forecasting. *Third*, instead of using clean human annotated ground-truth, our work uses noisy supervision from off-the-shelf methods (Section 3.1), transferring the focus from expensive problem specific annotations to robust learning algorithms.

3. Proposed Method

We frame the task of forecasting human locomotion in egocentric view (of the vehicle) as a sequence-to-sequence

problem. Suppose p_t denotes the human pose at time t , which contains d two-dimensional joints, *i.e.*, $p_t = \{(u^i, v^i)\}_{i=1}^d = \{\vec{u}_i\}_{i=1}^d$. Forecasting human locomotion is then defined as: given t_p previous poses, $\{p_i\}_{i=t-t_p+1}^t \equiv \{p_{t-t_p+1}, \dots, p_t\}$, predict the position of each of these joints for t_f future frames. In other words, predict the sequence $\{p_i\}_{i=t+1}^{t+t_f} \equiv \{p_{t+1}, \dots, p_{t+t_f}\}$. Our method is illustrated in Fig. 2. We start with noisy pose estimates generated using state of the art pretrained models (Section 3.1). The missing and low confidence joints are filled in using the pose completion module (Section 3.2). The completed poses are then split into global and local streams (Section 3.3) that are forecasted using Quasi Recurrent Neural Networks [34] (Section 3.4) in an Encoder-Recurrent-Decoder pose prediction module (Section 3.5). These predictions combine to forecast the future pedestrian locomotion.

3.1. Machine Generated Noisy Supervision

Keypoint Detection We use state-of-the-art models for *multiple-person keypoint detection* module to autonomously generate dense but noisy frame-level supervision for human poses p_t . To the best of our knowledge, there are no existing datasets with manually annotated pose data that can be used for the task of human locomotion forecasting in the wild. So, we use pre-trained models on task-specific datasets to infer pose and depth on our datasets. For each frame in the video, we create masked frames where everything except the pedestrians are masked out using the human labelled bounding box annotations. These masked frames are then processed through OpenPose [6, 7, 8], a pre-trained pose detection model to generate pose labels for every pedestrian in the frame. We notice that using masked full frames for gen-

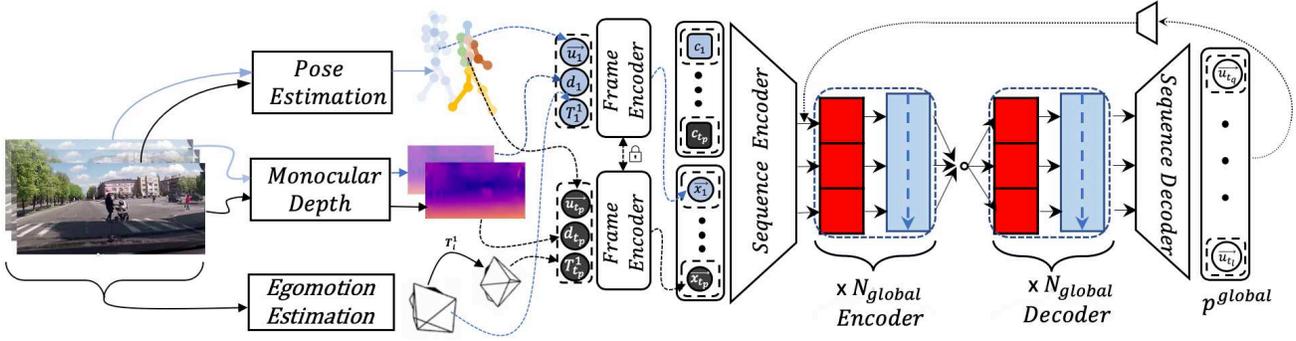


Figure 3: Encoder-Recurrent-Decoder architecture for forecasting p^{global} . \vec{u}_α , d_α and c_α represent coordinates, estimated monocular depth and confidence in joint \tilde{i} at frame α . \mathcal{T}_α^β represents the scene transformation matrix between frames β and α . The lock denotes the sharing of the frame encoder weights across different time steps of the input sequence. Dotted squares contain values concerned with the same frame.

erating frame-level pose supervision is faster and requires significantly lesser storage (upto 100 times lesser) than processing each cropped out pedestrian one by one without any observable decrease in pose detection accuracy. For the i^{th} joint detected in the t^{th} frame, OpenPose detects the 2D coordinates (u_t^i, v_t^i) and provides a confidence scores c_t^i , which is then used for denoising these detections as described in Section 3.2. These labelled keypoints form a human pose with $d = 25$ joints (Fig. 6).

Monocular Depth Estimation We autonomously estimate depth in a monocular camera using SuperDepth [35]. SuperDepth extends a subpixel convolutional layer for depth super-resolution, where high-resolution disparities are synthesized from their corresponding low-resolution convolutional features. We train the depth estimation model in a self supervised fashion as proposed in [35]. Fig. 4 shows qualitative examples of the depth estimation by SuperDepth. **Egomotion Estimation** We use the state-of-the-art-model unsupervised model [36] for autonomously estimating the camera motion that occurs between consecutive frames due to the movement of the egovehicle. However, since the method proposed in [36] can reliably estimate the scene transformation matrix \mathcal{T}_i^j between frames i and j only frame very close to each other ($\|i - j\| \leq 5$ in a 30 fps video) and the time horizons relevant for prediction are much larger (~ 30 frames), we chain the estimate between different overlapping short windows to obtain \mathcal{T}_{i+k}^i for $k > 5$.

3.2. Pose Completion

Off-the-shelf pre-trained pose detectors often result in noisy detection of keypoints. In addition, the small scale and off co-planar orientation of pedestrians in real-world 2D videos, the pose detector may miss detecting some keypoints even for a completely visible occlusion-free pedestrians in the frame. Hence, we propose a pose completion network for completing the detected human poses obtained from OpenPose. This processing has a two-fold benefit. *First*, it

fills in the joints that are not detected by the pose detection module and so tackles the missing data problem. It also suppresses noise by filling in the low confidence output with better estimates. *Second* and importantly, it enables us to decompose the motion (Section 3.3) with noisy data. Separating uncompleted global and local components of motion would be perplexing as the joints flicker frequently.

Network Architecture We use an over-complete symmetric auto encoder similar to [33, 32]. The network is only trained on the subset of the total training data that has been assigned high confidence scores by Openpose. In other words, the poses are filtered by thresholding on OpenPose confidence scores to only include the examples with $c_t^i > \alpha_c, \forall i \in \{1, \dots, d\}$. These high confidence examples are then used to train a symmetric autoencoder with dropout that embeds the $2d$ -dimensional vectors p_t to the latent dimension d_{ae} . Then, it maps them back to the denoised detections \hat{p}_t . Training with dropout in the input layer on high confidence examples, helps model the effect of missing data. Furthermore, supervising the loss function with the ground-truth on these good examples and the information bottleneck in the form of a narrow layer in the middle of network allows the model to learn to reconstruct the complete pose even if some of the joints are missing. The training procedure details are in Section 4.2.

Completion The trained autoencoder is then used to estimate the full pose on all the low confidence noisy detections

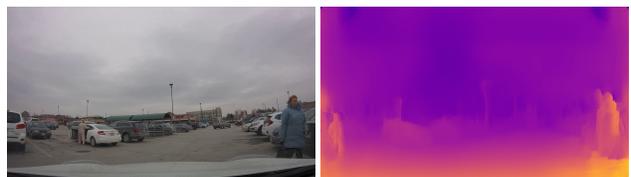


Figure 4: Examples of the input (left col.) and the output (right col.) of the monocular depth estimation module [35].

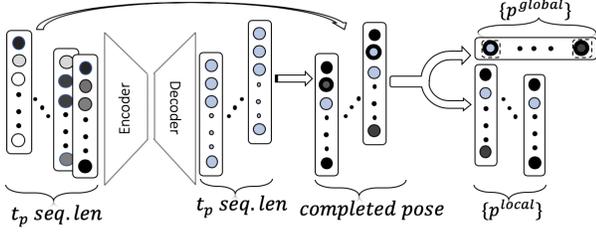


Figure 5: Pose completion and disentangling module. The shades represent the confidence in locating the joint (Black: highest confidence & white: missing data). Low confidence estimates are replaced with the autoencoder estimates (sky blue). It is then split into local & global streams.



Figure 6: Qualitative results showing the results of pose completion (Section 3.2). From left to right: denoised pose, cropped version of the same pedestrian with the original OpenPose output (green), denoised pose (blue).

from the pose detection module. These estimates then fill in the missing and low confidence joint detections. Mathematically, $(u_t^i, v_t^i) \leftarrow (\hat{u}_t^i, \hat{v}_t^i) \forall (i, t)$ such that $c_t^i \leq \alpha_c$.

3.3. Disentangling Pedestrian Locomotion into Global and Local Streams

After noise suppression using the autoencoder and filling in the missing joints, we propose to disentangle the overall pose locomotion sequence $\{p_j\}_{j=t-t_p+1}^t$ into two streams. A global motion stream encodes the overall rigid body motion of the pedestrian $\{p_j^{\text{global}}\}_{j=t-t_p+1}^t$ and a local motion stream encodes the motion of the body with respect to the global stream $\{p_j^{\text{local}}\}_{j=t-t_p+1}^t$. We hypothesize that these separate streams capture different granularities of simultaneous concurrent motion and need to be modelled and forecasted separately. The global motion stream p^{global} models large scale movements of the pedestrian position with respect to the camera, such as those caused by the trajectory or egocentric motion of the camera mounted on the moving car *etc.* The local motion stream captures the effect of depth change on the overall pose size and the movement of different joints of the pedestrian with respect to the global motion stream such small repetitive motions like swinging of arms.

Our proposal to disentangle global and local motion is motivated by the relative difference in the nature and mag-

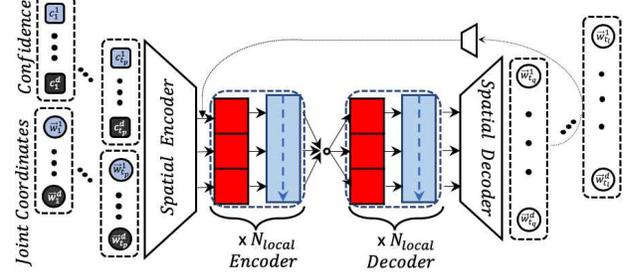


Figure 7: Architecture for forecasting the local stream, p^{local} . w_α^i represent the relative coordinates of joint i with respect to joint \tilde{i} at frame α derived from the pose completion module. Rest convention is same as Figure 3.

nitude of motion exhibited by these streams. Consider for example, the actual path $\{p_j\}_{j=t-t_p+1}^{t+t_f}$ traced by the elbow though space. In the original formulation this trajectory is quite complex, consisting of several swirls generated from gyrating around the shoulder and drifting with the rest of the body. In addition, this disentangling allows to significantly reduce the overall complexity, since each of the streams now model a much simpler and easier to predict motion.

We propose to use the neck joint sequence denoted by $(u_j^{\tilde{i}}, v_j^{\tilde{i}})_{j=t-t_p+1}^t$ as a representation of the global stream, because the neck is the most widely observed joint in the dataset. Note that this representation is possible only after filling in the missing data as discussed in Section 3.2. The local stream p^{local} (denoted by \vec{w}) is then derived by reparameterizing the original stream p_t (denoted by \vec{u}) as

$$\vec{w}_t^i = (u_t^i, v_t^i) - (u_t^{\tilde{i}}, v_t^{\tilde{i}})$$

$$\forall i \in \{1, \dots, d\}, i \neq \tilde{i}, t \in \{t - t_p + 1, \dots, t\}$$

3.4. Quasi-RNN Encoder-Decoder

The Quasi-Recurrent Neural Network [34] forms the backbone of the sequence-to-sequence learning structure described in Section 3.5. QRNNs consist of alternating convolutional and recurrent pooling module and is designed to parallelize efficiently better than vanilla LSTMs. Similar to the findings of Bradbury *et al.* [34], we found that Quasi-RNN trains faster (825 ms/batch) compared to LSTM (886 ms/batch) on a single GPU under same parameters and yield faster convergence for similar model capacity.

Network Architecture Our quasi-RNN encoder-decoder has N layers of alternate convolutions and recurrent pooling, both in the input encoder and the output decoder. As illustrated in Figs. 3 and 7, the recurrent pooling is a thin aggregation function applied to the convolutional activations. The encoder churns through the latent representation of the t_p previous poses and encodes the necessary information into a context vector. This vector is then consumed by the

QRNN decoder to forecast the future t_f poses mapped back to the same latent space. The dotted arrow represents the use of the technique of teacher forcing [37] for training recurrent networks for motion prediction, as also proposed by [2].

3.5. Forecasting the Future

After being split into separate granularities, the global and local streams are then forecasted separately with different prediction modules as described in the following.

Local Stream Forecasting As shown in Fig. 7, the filled in and decomposed pose $\{p_j^{\text{local}}\}_{j=t-t_p+1}^t$ is used as the input to the pose prediction module. This module comprises of a spatial encoder with the latent dimension d_{ae} . The weights of this spatial encoder are separately trained using the auto-encoder while the complexity of the latent space is similar. The forecasting is processed in the latent space with N_{local} layers of the QRNN Encoder Decoder module, briefly described in Section 3.4. We use the latent space to forecast because as reported in [32, 10] and also as confirmed by our pose completion module experiments, the human pose lies on a low dimensional manifold because of the various kinematic constraints enforced by the human body. Thus, forecasting in this lower dimensional denser space makes the prediction easier for the quasi RNN module. The predicted latent pose is then mapped back into the image space with the spatial decoder to forecast $\{p_j^{\text{local}}\}_{j=t+1}^{t+t_f}$.

Global Stream Forecasting As mentioned in Section 3.3, the global stream captures the effects of camera movement, perspective distortion from changing depth, and also the large-scale motion of target itself. Distortions due to the egocentric view, like perspective effects, can be implicitly learned using absolute coordinate trajectories while training. However, such an approach adds unnecessary complexity to the input distribution and hence is data inefficient. Instead, we propose to predict residuals from the first observed positions instead of forecasting absolute coordinates. In particular, we learn to predict the global stream from separately processed low level vision signals (monocular depth, camera egomotion). Furthermore, since these signals are machine generated and noisy, instead of using classical computer vision methods to employ these features like re-projecting the future positions with estimated camera pose matrices to cancel out egomotion effects, we propose a frame level encoder as depicted in Fig.3. The frame-level encoder is a feedforward neural network, which at sequence frame α takes as input observed coordinates $(\tilde{u}_{\alpha}^{\tilde{i}})$, estimated monocular depth d_{α} of joint \tilde{i} , and a flattened 3×4 transformation matrix T_{α}^1 estimating the change in the 6D pose between start of the sequence and frame α . The frame-level encoder architecture is motivated by the classical computer vision algorithm of projecting a point (p_1, p_2, p_3) to a new point (p'_1, p'_2, p'_3) under a transformation matrix \mathcal{T} and hence is shared across different time steps. The output of encoder is

a 2 dimensional vector \vec{x}_{α} , which is then pooled across time steps to be fed into the sequence level encoder. The rest of the architecture is similar to local stream forecasting module but with N_{global} layers of QRNN.

Loss Metric Both streams are trained with the ℓ_1 loss between the predicted pose and the original pose (before completion) weighted by the confidence scores $c_{i,j=t+1}^{j=t+t_f}$ of OpenPose in the original detection.

3.6. Recombining the Streams

These future predictions are finally recombined to forecast the overall human locomotion $\{p_j\}_{j=t}^{t+t_f}$ as follows:

$$\hat{p}_j = \begin{cases} p_j^{\text{global}} & \text{for } i = \tilde{i} \\ p_j^{\text{global}} + p_j^{\text{local},i} & \text{for } i \neq \tilde{i} \end{cases}$$

$\forall j \in \{t+1, \dots, t+t_f\}$ and $\forall i \in \{1, \dots, d\}$. This operation forms the inverse of the splitting procedure (Section 3.3).

4. Experiments

JAAD Dataset [38]: Joint Attention in Autonomous Driving is a dataset for modelling attention in the context of autonomous driving scenarios. The dataset contains 346 videos with 82032 frames, where the videos are the recorded with front-view camera under various scenes, weathers, and lighting conditions. It is also one of the very few autonomous driving datasets that has temporally dense pedestrian bounding box annotations. For example, popular egocentric datasets like KITTI [39], Berkeley DeepDrive [40], and CityScapes [41] do not contain human bounding boxes in a dense per-frame fashion. Similarly, PedX [42] and Mapillary [43] has pedestrian bounding box annotations, but these are separate images and not videos. The videos in JAAD range from short 5 second clips to longer 15 second videos shot through wide angle camera mounted inside the car in centre of windshield. It contains 64 clips of resolution 1920×1080 pixels and rest 293 of resolution 1280×720 pixels. Employing the procedure mentioned in 3.1, a little under 1500 non-overlapping 30 frame length human locomotion examples are extracted which are then split into train and test sets while taking care not to have the same pedestrian’s trajectories in both groups.

4.1. Experimental Settings

We use a $d = 25$ pose keypoints as extracted from OpenPose. Empirically, the value of α_c is set to 0.25 for all our following experiments. The pose completion network is trained with the ADAM optimizer [44] with learning rate 10^{-3} and dropout 0.5 applied to the input layer. The local stream network is trained with ADAM too with learning rate of 10^{-4} and $N_{\text{local}} = 4$. The global stream network is trained on the residuals between consecutive frames with ADAM at learning rate 0.001 and $N_{\text{global}} = 2$. We

Table 1: Average KDE for baselines and our proposed method. Despite their success at modelling other sequence learning tasks, their direct application in predicting human locomotion in egocentric view fail due to the complicated interaction between the local and global structures.

Method	Decomposition	KDE
Last Observed-Velocity	-	195.5
Constant-Velocity	-	48.9
Zero-Velocity	-	40.4
LSTM-ED [45]	✗	24.4
TCNN [46]	✗	31.8
GRU-ED [47]	✗	23.8
Structural RNN [3]	✗	31.7
Ours	✓	10.9

Table 2: Ablation Study on our method showing the relative importance of different components.

Stream	Completion	Disentanglement	KDE
Global	✗	-	6.8
Global	✓	-	5.6
Local	✗	-	9.5
Local	✓	-	6.3
Both	✗	✗	22.1
Both	✓	✗	18.2
Both	✗	✓	15.4
Both	✓	✓	10.9

use Keypoint Displacement Error (KDE) defined as the average ℓ_2 distance between the predicted and the ground truth keypoints to measure model’s performance, defined as $KDE = (\sum_{j=t+1}^{t+t_f} \|\hat{p}_j - p_j\|_1) / t_f$. Naturally, Mean KDE measures the Keypoint displacement error per human joint. For global stream, we use the same formulation but with $d = 1$ in contrast to $d = 25$ for local stream.

4.2. Pose Completion Results

As described in Section 3.2, we use a symmetrical over complete auto-encoder for pre-processing noisy detections. Our model comprises of a compressing module that encodes the $2d = 50$ dimensional pose into $d_{ae} = 10$ dimensions, in three equally proportionate non linear layers. Symmetrically, the decoder decompresses the latent pose to the original 50 dimensional space. A dropout of 0.5 dropout is applied to input layer and training is done only on poses with all keypoints having confidence more than c_t . This is trained as mentioned in Section 4.1 to yield 9.6 unweighed Mean KDE on high confidence pose detections on JAAD dataset. In practice, this means that on average the module is able to predict a missing joint within an ℓ_1 ball of radius 10 pixels centered at the joint’s true location in a frame. Considering that the average resolution of frames in JAAD is 1408×790

Table 3: KDE comparison of our method with GRU-ED (best performing baseline) for different output time horizons for the global stream with 15 frames input (half a second).

t_f	5	10	15	20	25	30
GRU-ED	5.7	6.6	9.7	11.4	14.1	18.3
Our Method	4.3	5.1	5.6	7.8	10.4	11.4

pixels and the pose completion network is not trained to regress on residuals but instead predicts the absolute coordinates in original frame resolution, this means that the error is about 1% – 1.5% of the frame dimensions and quite tolerable for our forecasting task.

4.3. Stream Forecasting

We compare our proposed method with several other baselines trained from scratch with the same input (*i.e.* un-completed pose estimates from Openpose) and output data. These methods are briefly described below:

Encoder Decoder Architectures : We train the popular Encoder Decoder architecture for sequence-to-sequence learning with several backbones like LSTM [48, 45], Gated Recurrent Units [47] and recently introduced Quasi RNN [34]. All the models are trained from scratch on the same input data and confidence values but without any explicit spatial encoding. Also, Quasi RNN is the most parallelizable and train much faster than other recurrent counterparts. Hence, it forms the backbone of our method.

Temporal Convolutional Neural Networks: Lea et al. [46] recently proposed Temporal Convolutional Networks which have also shown promising performance on a variety of sequence-to-sequence tasks. We found that TCNN trained with the same input signals, are the fastest to train and perform the best among deep learning methods without decomposition but are still much worse than the simple zero-velocity baseline. While TCNN perform well before decomposition, the future time horizon length t_f is constrained by the input sequence length t_p and does not allow flexible predictions as with other recurrent architectures.

Structural RNN[3]: SRNN has also achieved state of the art performance on several sequence prediction tasks including pose forecasting. However, we observe that its capacity is much larger than other sequence-to-sequence models because of larger number of trainable parameters that does not generalize well to our task (limited data regime).

Uniform Velocity: We also compare our method with simple yet powerful baselines [9, 49] based on zero, constant, and last observed velocity assumptions. In Zero Velocity, the future is forecasted assuming no movement from the last observed position. Similarly, in constant-velocity/last-observed-velocity, the forecasted velocity is assumed to be identically equal to the average observed velocity/last-



Figure 8: Qualitative pose prediction results. (A) the $t_p = 15$ length input sequence of poses. (B) and (C) the cropped frames with the pedestrian and the corresponding filled in pose at the start and end of the input sequence. (D) the prediction pedestrian locomotion for the next $t_f = 15$ frames. (D), (E), and (F) the predicted poses at the start, intermediate, and end of the output sequence, respectively. Video visualizations are available at: <https://karttikeya.github.io/publication/plf>

observed-velocity and this assumption is rolled out in time to predict future positions.

Results Discussion: Table 1 reports the Average KDE achieved by these approaches for predicting the pedestrian locomotion for $t_f = 15$ frames into the future given the previous $t_p = 15$ frames. Our method performs better than several widely used sequence-to-sequence models and also state-of-the-art pose forecasting models like Structural RNN [3]. Figure 8 shows some qualitative results from our method for pedestrian locomotion prediction.

Ablation: We perform the ablation study on our proposed method to investigate the importance of each component. As Table 2 indicates, both disentanglement and completion are quite central to our method. Also, the significant increase in performance conferred by decomposing the streams confirms our hypothesis that disentanglement is essential to locomotion forecasting and reduces the complication of overlapping simultaneous movements. Also noteworthy is the decrease in performance due decomposing without pose completion. This backs up our proposition that completion is essential to disentanglement and data is subject to severe method-induced noise without pose completion.

Time Horizon Experiments: We extensively test our novel egocentric global stream (trajectory) forecasting module for various time horizons against the best performing baseline from Table 2. Table 3 shows these trends for time horizons from 165 ms to 1 second. We observe that as the horizon increases and the trajectories become longer, our method outperforms the baseline by an increasing and larger margin. This is because our model accounts for the camera motion, which can accumulate over longer time horizons leading to a large error in egocentric motion agnostic methods.

5. Discussion

Taking into consideration the results discussed in the last section, we would like to revisit our original hypothesis and discuss final remarks that also highlight the new findings.

Disentanglement: We hypothesize that disentanglement of

human dynamics into overlapping but simpler streams helps overall prediction. This is corroborated with the results and ablation experiments discussed before. We further suspect that disentanglement performed in our proposed manner is useful in predicting motion dynamics of any object, cognizant or inanimate. The overall prediction difficulty of the spatially disentangled parts is lesser than that of the whole.

Finetuning Pose Detection: The state of the art pose detection network is pretrained on general purpose pose datasets and is not adapted to this setting. The pose completion network finetunes these detection conditioning on the information that the majority of human poses for pedestrian locomotion prediction are upright, as is the case when pedestrian is standing, walking or running. Thus, it implicitly also works as an adaptation step to use pose regularities specific to the setting of locomotion prediction. Additionally, it fills in the missing data because as discussed in Section 3.2, it is the existence of the keypoints and not so much their accuracy that helps the decomposition.

6. Conclusion

We propose a human locomotion forecasting pipeline that disentangles the complex pedestrian locomotion into simpler tasks of global and local motion prediction. We also present a pose completion module and an Encoder-Recurrent-Decoder style pose prediction network. The pose completion module disentangles the overlapping motion and fills in the missing detected joints. Furthermore, the pose prediction network forecasts these streams separately which are finally merged for the final prediction. We demonstrate that our method improves the overall performance & achieves state-of-the art short term human locomotion forecasting in egocentric view.

Acknowledgments Toyota Research Institute (TRI) provided funds to assist the authors with their research. This article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entities.

References

- [1] Hsu-kuang Chiu, Ehsan Adeli, Borui Wang, De-An Huang, and Juan Carlos Niebles. Action-agnostic human pose forecasting. *arXiv preprint arXiv:1810.09676*, 2018.
- [2] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2891–2900, 2017.
- [3] Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5308–5317, 2016.
- [4] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016.
- [5] Xin Huang, Stephen McGill, Brian C Williams, Luke Fletcher, and Guy Rosman. Uncertainty-aware driver trajectory prediction at urban intersections. *arXiv preprint arXiv:1901.05105*, 2019.
- [6] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017.
- [7] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *CVPR*, 2016.
- [8] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. In *arXiv preprint arXiv:1812.08008*, 2018.
- [9] Wei-Chiu Ma, De-An Huang, Namhoon Lee, and Kris M Kitani. Forecasting interactive dynamics of pedestrians with fictitious play. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 774–782, 2017.
- [10] Isinsu Katircioglu, Bugra Tekin, Mathieu Salzmann, Vincent Lepetit, and Pascal Fua. Learning latent representations of 3d human pose with deep neural networks. *International Journal of Computer Vision*, 126(12):1326–1341, 2018.
- [11] Anirudh Vemula, Katharina Muelling, and Jean Oh. Social attention: Modeling attention in human crowds. In *Proceedings of the International Conference on Robotics and Automation (ICRA) 2018*, May 2018.
- [12] Julian FP Kooij, Fabian Flohr, Ewoud AI Pool, and Dariu M Gavrilă. Context-based path prediction for targets with switching dynamics. *International Journal of Computer Vision*, 127(3):239–262, 2019.
- [13] Julian FP Kooij, Nicolas Schneider, and Dariu M Gavrilă. Analysis of pedestrian dynamics from a vehicle perspective. In *2014 IEEE Intelligent Vehicles Symposium Proceedings*, pages 1445–1450. IEEE, 2014.
- [14] Takuma Yagi, Karttikeya Mangalam, Ryo Yonetani, and Yoichi Sato. Future person localization in first-person videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7593–7602, 2018.
- [15] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social gan: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2255–2264, 2018.
- [16] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, and Silvio Savarese. Sophie: An attentive gan for predicting paths compliant to social and physical constraints. *arXiv preprint arXiv:1806.01482*, 2018.
- [17] Haosheng Zou, Hang Su, Shihong Song, and Jun Zhu. Understanding human behaviors in crowds by imitating the decision-making process. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [18] A. Vemula, K. Muelling, and J. Oh. Social Attention: Modeling Attention in Human Crowds. *ArXiv e-prints*, October 2017.
- [19] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B Choy, Philip HS Torr, and Manmohan Chandraker. Desire: Distant future prediction in dynamic scenes with interacting agents. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 336–345, 2017.
- [20] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [21] Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5):4282, 1995.
- [22] D. Xie, S. Todorovic, and S. Zhu. Inferring dark matter and dark energy from videos. In *2013 IEEE International Conference on Computer Vision*, pages 2224–2231, Dec 2013.
- [23] David Ellis, Eric Sommerlade, and Ian Reid. Modelling pedestrian trajectory patterns with gaussian processes. In *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*, pages 1229–1234. IEEE, 2009.
- [24] Jürgen Wiest, Matthias Höffken, Ulrich Kreßel, and Klaus Dietmayer. Probabilistic trajectory prediction with gaussian mixture models. In *2012 IEEE Intelligent Vehicles Symposium*, pages 141–146. IEEE, 2012.
- [25] Yu Yao, Mingze Xu, Chiho Choi, David J Crandall, Ella M Atkins, and Behzad Dariush. Egocentric vision-based future vehicle localization for intelligent driving assistance systems. *arXiv preprint arXiv:1809.07408*, 2018.
- [26] Chen Li, Zhen Zhang, Wee Sun Lee, and Gim Hee Lee. Convolutional sequence to sequence model for human dynamics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5226–5234, 2018.
- [27] Liang-Yan Gui, Yu-Xiong Wang, Xiaodan Liang, and José MF Moura. Adversarial geometry-aware human motion prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 786–803, 2018.

- [28] Partha Ghosh, Jie Song, Emre Aksan, and Otmar Hilliges. Learning human motion models for long-term predictions. In *2017 International Conference on 3D Vision (3DV)*, pages 458–466. IEEE, 2017.
- [29] Mengyao Zhai, Jiacheng Chen, Ruizhi Deng, Lei Chen, Ligeng Zhu, and Greg Mori. Learning to forecast videos of human activity with multi-granularity models and adaptive rendering. *arXiv preprint arXiv:1712.01955*, 2017.
- [30] JY Chang. Dr-net: denoising and reconstruction network for 3d human pose estimation from monocular rgb videos. *Electronics Letters*, 54(2):70–72, 2017.
- [31] Dana H Ballard. Modular learning in neural networks. In *AAAI*, pages 279–284, 1987.
- [32] Bugra Tekin, Isinsu Katircioglu, Mathieu Salzmann, Vincent Lepetit, and Pascal Fua. Structured prediction of 3d human pose with deep neural networks. *arXiv preprint arXiv:1605.05180*, 2016.
- [33] Matthew Trumble, Andrew Gilbert, Adrian Hilton, and John Collomosse. Deep autoencoder for combined human pose estimation and body model upscaling. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 784–800, 2018.
- [34] James Bradbury, Stephen Merity, Caiming Xiong, and Richard Socher. Quasi-recurrent neural networks. *arXiv preprint arXiv:1611.01576*, 2016.
- [35] Sudeep Pillai, Rares Ambrus, and Adrien Gaidon. Superdepth: Self-supervised, super-resolved monocular depth estimation. *arXiv preprint arXiv:1810.01849*, 2018.
- [36] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, 2017.
- [37] Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280, 1989.
- [38] Iuliia Kotseruba, Amir Rasouli, and John K Tsotsos. Joint attention in autonomous driving (jaad). *arXiv preprint arXiv:1609.04741*, 2016.
- [39] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361. IEEE, 2012.
- [40] Fisher Yu, Wenqi Xian, Yingying Chen, Fangchen Liu, Mike Liao, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving video database with scalable annotation tooling. *arXiv preprint arXiv:1805.04687*, 2018.
- [41] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [42] Wonhui Kim, Manikandasriram Srinivasan Ramanagopal, Charles Barto, Ming-Yuan Yu, Karl Rosaen, Nick Goumas, Ram Vasudevan, and Matthew Johnson-Roberson. Pedx: Benchmark dataset for metric 3d pose estimation of pedestrians in complex urban intersections. *IEEE Robotics and Automation Letters*, 2019.
- [43] Gerhard Neuhold, Tobias Ollmann, Samuel Rota Bulò, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *International Conference on Computer Vision (ICCV)*, 2017.
- [44] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [45] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
- [46] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. Temporal convolutional networks for action segmentation and detection. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 156–165, 2017.
- [47] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [48] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [49] Christoph Schöller, Vincent Aravantinos, Florian Lay, and Alois Knoll. The simpler the better: Constant velocity for pedestrian motion prediction. *arXiv preprint arXiv:1903.07933*, 2019.