

Uncertainty-aware State Space Transformer for Egocentric 3D Hand Trajectory Forecasting

Wentao Bao^{1*}, Lele Chen³, Libing Zeng⁴, Zhong Li², Yi Xu², Junsong Yuan⁵, Yu Kong¹

¹Michigan State University, ²OPPO US Research Center, ³Sony AI,

⁴Texas A&M University, ⁵University at Buffalo

{baowenta,yukong}@msu.edu, lele.chen@sony.com, libingzeng@tamu.edu,

{zhong.li,yi.xu}@oppo.com, jsyuan@buffalo.edu

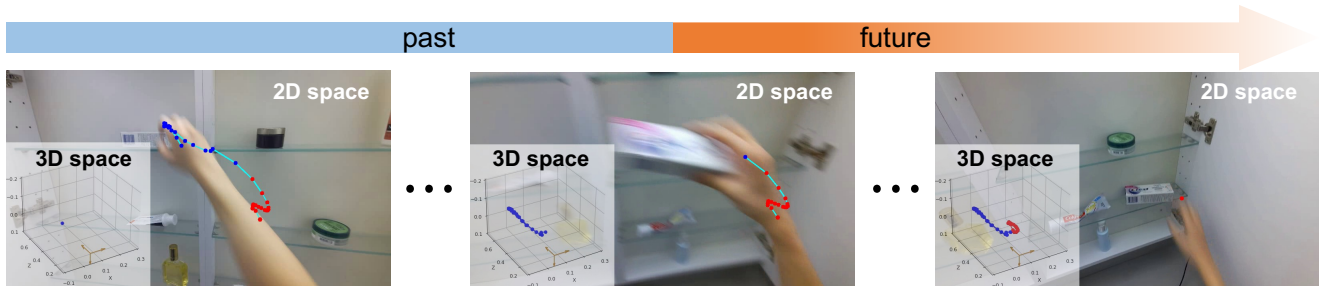


Figure 1: **Egocentric 3D Hand Trajectory Forecasting.** Our goal is to predict the future 3D hand trajectory (in red) given the past observations of egocentric video and trajectory (in blue). Compared to the 2D image space, predicting trajectory in a global 3D space is practically more valuable to understand human intention and behavior in AR/VR applications.

Abstract

Hand trajectory forecasting from egocentric views is crucial for enabling a prompt understanding of human intentions when interacting with AR/VR systems. However, existing methods handle this problem in a 2D image space which is inadequate for 3D real-world applications. In this paper, we set up an egocentric 3D hand trajectory forecasting task that aims to predict hand trajectories in a 3D space from early observed RGB videos in a first-person view. To fulfill this goal, we propose an uncertainty-aware state space Transformer (USST) that takes the merits of the attention mechanism and aleatoric uncertainty within the framework of the classical state-space model. The model can be further enhanced by the velocity constraint and visual prompt tuning (VPT) on large vision transformers. Moreover, we develop an annotation workflow to collect 3D hand trajectories with high quality. Experimental results on H2O and EgoPAT3D datasets demonstrate the superiority of USST for both 2D and 3D trajectory forecasting. The code and datasets are publicly released: <https://github.com/Cogito2012/USST>.

*This work was done when Wentao Bao and Libing Zeng were interns mentored by Lele Chen at OPPO US Research Center.

1. Introduction

Egocentric video understanding aims to understand the camera wearers’ behavior from the first-person view. It is receiving increasing attention in recent years [18, 27, 16, 31, 29, 36, 46, 32, 8, 7, 28] due to its analogousness to the way human visually perceives the world. An important egocentric vision task is to forecast the egocentric hand trajectory of the camera wearer [34], which has great value in Augmented/Virtual Reality (AR/VR) applications. For example, the predicted 3D trajectories can help plan and stabilize a patient’s 3D hand motion who has upper-limb neuromuscular disease [27]. Besides, the early predicted 3D hand trajectory is key to reducing rendering latency in VR games for achieving an immersive gaming experience [15].

In the existing literature, egocentric 3D hand trajectory forecasting is far from being explored. The method in [34] could only predict 2D trajectory on an image and cannot forecast precise 3D hand movements. Recent works [5, 45, 49] predict the trajectory or 3D human motions from egocentric views, but they do not predict the 3D trajectory of the camera wearer. Besides, though forecasting the 3D hand pose provides fine-grained information about 3D hands [10], it is out of our scope as we focus on the camera wearers’ planning behavior revealed by 3D hand trajectory.

The challenges of egocentric video-based 3D trajectory forecasting are significant. First, accurate large-scale 3D trajectory annotations are labor-intensive and expensive. They rely on wearable markers or multi-camera systems for hand motion capture in a controlled environment. Second, learning the depth of 3D trajectory from egocentric videos is challenging. On one hand, using 2D video frames to estimate 3D trajectory depth is an ill-conditioned problem similar to other monocular 3D tasks [42, 38, 2]. Even if the historical 3D hand trajectory is utilized as the input, how to exploit the visual and trajectory information for forecasting is still nontrivial. On the other hand, due to the inevitable camera motion in an egocentric view, the background of the scene is visually dynamic which poses a significant barrier to inferring the foreground depth [30, 52]. Third, as a Seq2Seq forecasting problem, it is critical to formulate the latent transition dynamics [17, 3] that allows the variances of data due to anytime forecasting and limited observations.

In this paper, we address these challenges by developing an uncertainty-aware state space transformer (USST). It follows the state-space model [40] by taking the observed egocentric RGB videos with the historical 3D trajectory as input to predict future 3D trajectory. Our model deals with the depth noise of trajectory annotation by introducing the depth robust aleatoric uncertainty in training. To fuse the information from the dense RGB pixels and sparse historical trajectory, we leverage visual and temporal transformer encoders as backbones and utilize the recent visual prompt tuning (VPT) to enhance the visual features. Following the state space model, we develop a novel attention-based state transition module and an emission module with a predictive link to predict the 3D global trajectory coordinates. Moreover, to take the hand motion inertia into consideration, we propose a velocity constraint to regularize the model training, which helps generalize to unseen scenarios.

To enable egocentric 3D hand trajectory prediction, we follow [27] to develop a scalable annotation workflow to automate the annotation on RGB-D data from head-mounted Kinect devices. In particular, camera motion is estimated to transform the 3D trajectory annotations from local to global camera coordinate system. Experimental results on H2O and EgoPAT3D datasets show that our method is effective and superior to existing Transformer-based approaches [34] and other general Seq2Seq models. In summary, our contributions are as follows:

- We propose an uncertainty-aware state space transformer (USST) that consists of a novel state transition and emission, aleatoric uncertainty, and visual prompt tuning, which are empirically found effective.
- We collected and will release our annotations on H2O [25] and EgoPAT3D [27] datasets that will benefit the egocentric 3D hand trajectory forecasting research.

- We benchmarked recent methods on the proposed task and experimental results show that our method achieves the best performance and could be generalizable to unseen egocentric scenes.

2. Related Work

Trajectory Prediction Predicting the physical trajectory of moving objects is a long-standing research topic. It has been widely studied in applications for pedestrians [1, 51], vehicle [9, 35]. Many of them are developed for the third-person view and predict trajectory in 2D pixel space. Given that the first-person view is more realistic in AR/VR applications, recent few works [34, 33] are trying to predict the hand-object interaction from egocentric videos. Though the method in [5] predicts the pedestrian trajectory in 3D space, their method primarily addresses the social interaction of multiple pedestrians. The recent work [39] also targets the egocentric 3D trajectory for pedestrian scenarios, but their method leverages depth modality and nearby person’s trajectory as context information which are practically uneasy to collect. Besides, due to the annotation noise and uncertain nature of trajectory prediction, probabilistic modeling is widely adopted in existing literature [47, 43, 35]. In this paper, following the probabilistic setting, we step toward the egocentric 3D hand trajectory prediction using practically accessible RGB videos for AR/VR scenarios.

Egocentric Video Representation Egocentric videos are recorded in a first-person view. Different from the videos in a third-person view, learning an egocentric video representation is more challenging due to the dynamic background caused by camera motion and implicit intention of activities from camera wearers [29, 12, 37]. For 3D trajectory prediction, existing commonly used egocentric video datasets such as EPIC-Kitchens [7, 8] do not provide the depth information and camera parameters, which are essential for annotating the 3D hand trajectory. Though the recent Ego4D [18] benchmark provides a hand forecasting subset, the annotations are defined as 2D locations in image space. Therefore, we resort to a cost-effective workflow to collect 3D hand trajectory annotations from existing 3D hand pose datasets such as the EgoPAT3D [27] and H2O [25] datasets. Our annotation workflow can be deployed to any egocentric dataset collected by head-mounted RGB-D sensors.

State Space Model State-space Model (SSM) originates from the control engineering field. It is conceptually general and inspires many classical SSMs such as the Kalman Filtering for prediction tasks. Recent deep SSMs [4, 6, 13] are increasingly popular by combining Recurrent Neural Networks (RNNs) with the Variational AutoEncoders (VAE). However, these approaches are limited in practice due to the complex long-term dependency on highly-structured se-

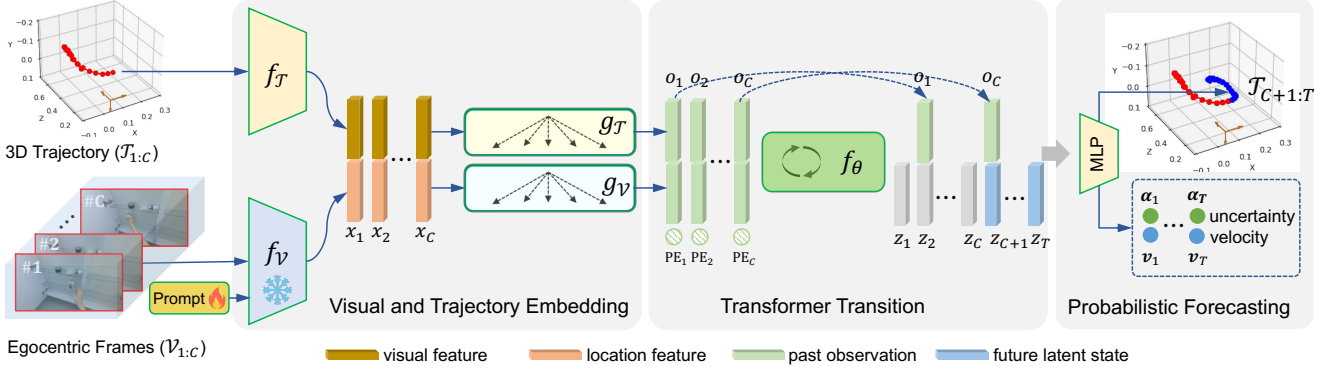


Figure 2: **Proposed USST Model.** Given the RGB frames and 3D hand locations of C observed time steps, we extract and concatenate their features as $\mathbf{x}_{1:C}$ by the prompted backbone f_V and MLP model f_T , which are further fed into transformer encoders to produce temporal observations $\mathbf{o}_{1:C}$. Together with positional encodings $\text{PE}_{1:T}$ of the full horizon, our state transition layer could recursively extrapolate the latent states $\mathbf{z}_{C+1:T}$ for 3D trajectory forecasting along with uncertainty and velocity (α and \mathbf{v}) in $T - C$ future steps.

quential data. In [40], a deep SSM is proposed by combining Kalman Filtering with deep neural networks. However, the linear Gaussian assumption is uneasy to hold for high-dimensional data in the real world. To address these challenges, ProTran [41] introduces a probabilistic Transformer [44] under a variational SSM for time-series prediction. However, maximizing the variational low bound of ProTran suffers from notorious KL vanishing issue [14]. AgentFormer [50] can also be regarded as a Transformer-based SSM, but its autoregressive decoding limits its efficiency to low-dimensional motion data. In this paper, we propose a Transformer-based SSM that allows for long-term dependency and latent dynamics efficiently in practice.

3. Proposed Method

Problem Setup As shown in Fig. 1, the Egocentric 3D hand trajectory forecasting model takes as input C observed RGB frames $\mathcal{V}_{1:C} = \{\mathbf{I}_1, \dots, \mathbf{I}_C\}$ and 3D hand trajectory $\mathcal{T}_{1:C} = \{\mathbf{p}_1, \dots, \mathbf{p}_C\}$ to predict the future 3D hand trajectory $\mathcal{T}_{C+1:T} = \{\mathbf{p}_{C+1}, \dots, \mathbf{p}_T\}$ in a finite horizon T . Here, $\mathbf{I}_t \in \mathbb{R}^{H \times W \times 3}$ and $\mathbf{p}_t = [x_t, y_t, z_t]^\top$ are egocentric RGB frame and 3D hand trajectory point at time step t , respectively. In practice, the 3D point \mathbf{p}_t is defined in a global 3D world coordinate system. The ultimate goal is to learn a model Φ by maximizing the expectation of the likelihood over the training dataset \mathcal{D} :

$$\max_{\Phi} \mathbb{E}_{\mathcal{V}, \mathcal{T} \sim \mathcal{D}} [p_{\Phi}(\mathcal{T}_{C+1:T} | \mathcal{T}_{1:C}, \mathcal{V}_{1:C})]. \quad (1)$$

In this paper, we formulate the problem as a state-space model. In the following sections, we will introduce the proposed model in detail.

3.1. Uncertainty-aware State Space Transformer

Existing SSMs could formulate the probabilistic nature of trajectory prediction. However, they do not explicitly

handle the data noise issue which is commonly encountered when using RGB-D sensors to get 3D trajectory annotations. To mitigate the uncertainty from data labeling, we follow the line of research [22, 20] and propose an uncertainty-aware state space transformer (USST) to handle the dynamics of 3D hand trajectory in egocentric scenes.

Formally, following the latent variable modeling, the probability in Eq. (1) over a full sequence $\mathcal{T}_{1:T}$ can be factorized by introducing T latent variables $\mathcal{Z}_{1:T}$:

$$p(\mathcal{T}_{1:T} | \mathcal{T}_{1:C}, \mathcal{V}_{1:C}) = \int p(\mathcal{T}_{1:T} | \mathcal{Z}_{1:T}) p(\mathcal{Z}_{1:T} | \mathcal{T}_{1:C}, \mathcal{V}_{1:C}) d\mathcal{Z}, \quad (2)$$

where the *state transition* $p(\mathcal{Z}_{1:T} | \mathcal{T}_{1:C}, \mathcal{V}_{1:C})$ and the *emission* $p(\mathcal{T}_{1:T} | \mathcal{Z}_{1:T})$ are learned from data \mathcal{D} . Following the SSM formulation, we propose to factorize the two terms by the independency assumptions:

$$p_{\theta}(\mathcal{Z}_{1:T} | \mathcal{T}_{1:C}, \mathcal{V}_{1:C}) = \prod_{t=1}^T p_{\theta}(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathbf{p}_{1:C}, \mathbf{I}_{1:C}), \quad (3)$$

$$p_{\phi}(\mathcal{T}_{1:T} | \mathcal{Z}_{1:T}) = \prod_{t=1}^T p_{\phi}(\mathbf{p}_t | \mathbf{z}_t, \mathbf{p}_{t-1}),$$

where the latent variable $\mathbf{z}_t \in \mathcal{Z}_{1:T}$ is generated by taking as input \mathbf{z}_{t-1} and the previous trajectory point \mathbf{p}_t and RGB frame \mathbf{I}_t . To address the label noise issue, we formulate the emission model $p_{\phi}(\mathbf{p}_t | \mathbf{z}_t, \mathbf{p}_{t-1})$ as a probabilistic module to learn the aleatoric uncertainty. In the following paragraphs, we will elaborate on feature embedding, state transition, and probabilistic prediction.

Visual and Trajectory Embedding As advocated by [41, 50], Transformers are effective to capture the long-term dependency for sequential data. Therefore, we propose to leverage visual and temporal Transformers [44] as encoders to learn the features from the dense RGB frames and sparse

trajectory points. Specifically, we first embed the observed sequence of egocentric RGB frames and 3D trajectory by models f_V and f_T , followed by modality-specific transformers g_V and g_T . This process can be expressed as

$$\begin{aligned} [\mathbf{x}_t^{(V)}, \mathbf{x}_t^{(T)}] &= [f_V(\mathbf{I}_t), f_T(\mathbf{p}_t)], \\ \mathbf{o}_1^{(V)}, \dots, \mathbf{o}_C^{(V)} &= g_V(\mathbf{x}_1^{(V)}, \dots, \mathbf{x}_C^{(V)}), \\ \mathbf{o}_1^{(T)}, \dots, \mathbf{o}_C^{(T)} &= g_T(\mathbf{x}_1^{(T)}, \dots, \mathbf{x}_C^{(T)}), \end{aligned} \quad (4)$$

where f_V is a vision backbone, e.g., ResNet [19] and ViT [11]. f_T is implemented as MLP following [34, 27]. g_V and g_T are transformer encoders that consist of B stacked multi-head attention blocks. For each block b , a single-head attention block can be expressed as

$$\text{Attn}(\mathbf{Q}_b, \mathbf{K}_b, \mathbf{V}_b) = \text{softmax}\left(\frac{\mathbf{Q}_b \mathbf{K}_b^T}{\sqrt{d}} + \mathbf{M}\right) \mathbf{V}_b, \quad (5)$$

where $\mathbf{Q}_b, \mathbf{K}_b, \mathbf{V}_b \in \mathbb{R}^{T \times d}$ are the projected query, key, and value matrices from the output of the previous block $b-1$, i.e., $[\mathbf{Q}_b; \mathbf{K}_b; \mathbf{V}_b] = [\mathbf{W}_b^Q \mathbf{Q}_{b-1}; \mathbf{W}_b^K \mathbf{K}_{b-1}; \mathbf{W}_b^V \mathbf{V}_{b-1}]$. All the \mathbf{W}_b are learnable parameters. The binary mask $\mathbf{M} \in \mathbb{R}^{T \times T}$ zeros out the last $T-C$ columns and rows for the trajectory prediction problem. To capture global temporal interaction, the input of the first block $\mathbf{Q}_0, \mathbf{K}_0$, and \mathbf{V}_0 are the same as $\mathbf{x}_t + \text{PE}(t)$ where $\text{PE}(t)$ is the positional encoding for $t \in [1, T]$.

Transformer Transition With the encoded observations $\mathbf{o}_t = [\mathbf{o}_t^{(V)}; \mathbf{o}_t^{(T)}]$ where we use $[\cdot]$ to represent the feature concatenation, it is critical to formulate the state transition and future trajectory prediction based on Eq. (3). Inspired by [41], we propose an attention-based autoregressive module to formulate posterior $p_\theta(\mathbf{z}_t | \mathbf{z}_{1:t-1}, \mathbf{o}_{1:C})$. Specifically, we first embed \mathbf{o}_t with positional encoding by $\mathbf{h}_t = \text{LayerNorm}(\text{MLP}(\mathbf{o}_t) + \text{PE}(t))$. Then, the latent feature \mathbf{z}_t is recursively encapsulated by the hidden variables $\bar{\mathbf{w}}_t$ and $\hat{\mathbf{w}}_t$ by attention modules (illustrated in Fig. 3):

$$\begin{aligned} \bar{\mathbf{w}}_t &= \text{LayerNorm}([\mathbf{z}_{t-1}; \text{Attn}(\mathbf{z}_{t-1}, \mathbf{z}_{1:t-1}, \mathbf{z}_{1:t-1})]), \\ \hat{\mathbf{w}}_t &= \text{LayerNorm}([\bar{\mathbf{w}}_t; \text{Attn}(\bar{\mathbf{w}}_t, \mathbf{h}_{1:C}, \mathbf{h}_{1:C})]), \\ \mathbf{z}_t &= \text{LayerNorm}(\text{MLP}([\hat{\mathbf{w}}_t; \text{MLP}(\hat{\mathbf{w}}_t)]) + \text{PE}(t)), \end{aligned} \quad (6)$$

where the two multi-head attention modules capture the previously generated state $\mathbf{z}_{1:t-1}$ and the hidden states of all observations $\mathbf{h}_{1:C}$. Contrary to ProTran, we use concatenation $[\cdot]$ rather than addition before layer normalization. The insight behind this is that the queried feature from a historical context can be better preserved without being dominated by \mathbf{z}_{t-1} in addition operation. Moreover, we remove the stochasticity of \mathbf{z}_t and instead use a probabilistic decoder as introduced next to handle the dynamics of trajectory prediction. The benefit is that we avoid the KL divergence vanishing issue from optimizing the ELBO objective which is known to exist in variational recurrent models [14].

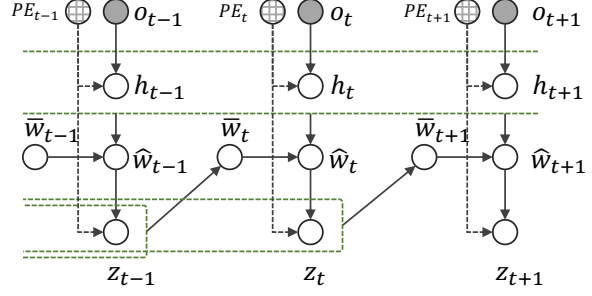


Figure 3: **Unrolled illustration of Eq. (6).** Bold arrows are learnable, and green dashed lines show the attention ranges.

Probabilistic Forecasting Instead of placing the stochasticity in the transition model $p_\theta(\mathbf{z}_t | \mathbf{z}_{1:t-1}, \mathbf{o}_{1:C})$, we propose to formulate the emission process $p_\phi(\mathbf{p}_t | \mathbf{z}_t, \mathbf{p}_{t-1})$ as a probabilistic model by predicting both the mean $\hat{\mathbf{p}}_t$ and variance $\hat{\sigma}_t^2$ of each 3D hand trajectory point:

$$[\hat{\mathbf{p}}_t, \hat{\alpha}_t] = [\text{MLP}([\mathbf{z}_t; \mathbf{o}_{t-1}^{(T)}]), \text{softplus}(\text{MLP}([\mathbf{z}_t; \mathbf{o}_{t-1}^{(T)}]))], \quad (7)$$

where the uncertainty $\hat{\alpha}_t := \log \hat{\sigma}_t^2$ to enable numerical stability. The trajectory point \mathbf{p}_t follows a predictive Gaussian distribution, i.e., $\mathbf{p}_t \sim \mathcal{N}(\hat{\mathbf{p}}_t, \hat{\sigma}_t)$. As the $\mathbf{o}_{t-1}^{(T)}$ encodes the observation from \mathbf{p}_{t-1} and its global historical context, our emission model is thus more powerful to predict \mathbf{p}_t . This predictive mode has also been successful in traditional methods such as the SRNN [13] and VRNN [6].

Discussion Compared to ProTran [41], our formulation could individually formulate the trajectory and visual context $\mathbf{p}_{1:C}$ and $\mathbf{I}_{1:C}$ in state transition by modality-specific embeddings and the predictive link from \mathbf{p}_{t-1} to \mathbf{p}_t , while ProTran only handles single modality context $\mathbf{p}_{1:C}$ in state transition and emission. In addition, to learn model parameters θ and ϕ , ProTran has to use variational posterior distribution $q_\theta(\mathbf{z}_t | \mathbf{z}_{1:t-1}, \mathbf{p}_{1:C})$ to help approximate Eq. (2) by ELBO maximization. In contrast, our method does not need approximation and formulates the data uncertainty of \mathbf{p}_t from a Bayesian perspective (Eq. (6)), which is empirically more effective to handle data noise.

3.2. Model Training

Depth Robust Aleatoric Uncertainty With the predicted trajectory points $\hat{\mathbf{p}}_t$ along with the uncertainty $\hat{\alpha}$, according to [22, 3], the model training is essentially to learn the heteroscedastic aleatoric uncertainty (HAU) from data. As shown in [23, 20], by minimizing the KL divergence between the predictive Gaussian distributions and Dirac distribution of the ground truth trajectory, the objective in Eq. (1)

is equivalent to minimizing the HAU loss at each time t :

$$\mathcal{L}_{\text{haul}}(\hat{\alpha}, \hat{\mathbf{p}}, \mathbf{p}) = e^{-\hat{\alpha}} \sum_{i=1}^{|\mathbf{p}|} \|p_i - \hat{p}_i\|^2 + \hat{\alpha}, \quad (8)$$

where p_i and \hat{p}_i are 3D coordinate values (x, y, z) of ground truth and model prediction, respectively.

In our task, the trajectory depth z is more challenging to predict than x and y due to 1) the weak implicit correspondence between the past visual context $\mathcal{V}_{1:C}$ and future hand depth, and 2) more importantly, the inevitable annotation noise from depth sensors. To handle these challenges, we propose to decouple the aleatoric uncertainty into $\hat{\alpha}_t$ which is isotropic for (x, y) and $\hat{\beta}_t$ specifically for z , respectively. Then, the predictions of (x, y) and z are weighted by factors w_t so that the regression loss becomes

$$\mathcal{L}_{\text{DRAU}}^{(t)}(\hat{\mathbf{y}}, \hat{\alpha}) = \mathcal{L}_{\text{haul}}(\hat{\alpha}_t, \hat{\mathbf{p}}_t^{(2d)}, \mathbf{p}_t^{(2d)}) + w_t \mathcal{L}_{\text{haul}}(\hat{\beta}_t, \hat{z}_t, z_t), \quad (9)$$

where the weight w_t is determined by the negative temporal difference of ground truth $z_{1:T}$ with softmax normalization:

$$w_t = \frac{\exp(-\Delta z_t)}{\sum_{t=1}^T \exp(-\Delta z_t)}, \quad \Delta z_t = |z_t - z_{t-1}|. \quad (10)$$

Since Δz_t indicates the stability of depth transition, the motivation of w_t is to encourage large weight on the stable depth transitions (small Δz_t) in a trajectory, which enables the training to focus less on the unstable depth so that the model is robust to noisy depth annotations.

Velocity Constraints To explicitly inject the physical rule of hand motion into the model, we additionally take the motion inertia into consideration. Specifically, we leverage the transitioned states $\{\mathbf{z}_1, \dots, \mathbf{z}_T\}$ learned from Eq. (6) to predict the velocities $\{\mathbf{v}_1, \dots, \mathbf{v}_T\}$ by an MLP. Then, we propose the following velocity constraint in training:

$$\mathcal{L}_{\text{velo}}(\hat{\mathbf{v}}, \mathbf{p}) = \sum_{t=1}^T (\|\mathbf{p}_t - \mathbf{p}_{t-1} - \hat{\mathbf{v}}_t\|^2) + \gamma \sum_{t=C+1}^T \left(\|\mathbf{p}_C + \sum_{i=C+1}^t \hat{\mathbf{v}}_i - \hat{\mathbf{p}}_t\|^2 \right), \quad (11)$$

where the first term uses the first-order difference of locations \mathbf{p}_t to supervise the predicted velocity $\hat{\mathbf{v}}$ and we set \mathbf{p}_0 to zero. The second term is to constrain the future predicted trajectory point $\hat{\mathbf{p}}_t$ with the warped point, which is computed by adding the accumulative predicted velocities onto the last observed point \mathbf{p}_C since the time interval is one. Empirically, we found the velocity constraint enables better generalization capability to unseen data (see Table 3).

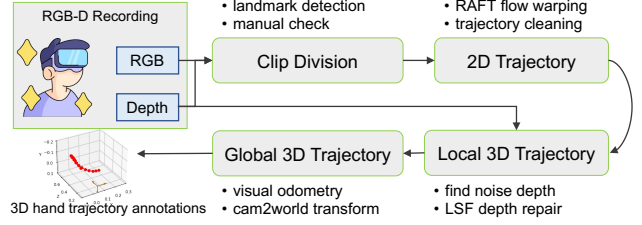


Figure 4: **Annotation Workflow.** For more details of the annotation procedure, please refer to our supplementary materials.

Visual Prompt Tuning Visual prompt tuning (VPT) [21] has been recently successful to adapt large visual foundation models to downstream vision tasks. In this paper, we leverage VPT to adapt pre-trained visual backbone $f_{\mathcal{V}}$ for the trajectory prediction task. The basic idea is to append learnable prompt embeddings \mathbf{P} to the input image \mathbf{I} and only learn a few layers of MLP head h_{ψ} while keeping the backbone parameters Ψ frozen as Ψ^* in training:

$$\mathbf{x}_t^{(\mathcal{V})} = h_{\psi}(f_{\mathcal{V}}^{\Psi^*}(\mathbf{I}_t, \mathbf{P})) \quad (12)$$

where only the head parameters ψ and visual prompt \mathbf{P} are learned in training. Since $\{\psi, \mathbf{P}\}$ are much smaller than Ψ , the VPT is highly efficient in training. We implemented $f_{\mathcal{V}}$ with both ResNet [19] and ViT [11], without noting significant performance difference. However, applying VPT achieves better 3D hand trajectory prediction performance than traditional fine-tuning (see Fig. 8). This is interesting as there is no existing literature that explores the VPT for vision-based trajectory prediction problems.

4. Experiments

4.1. Datasets

Since there is no available egocentric 3D hand trajectory dataset, we collect annotations based on two existing datasets, i.e., H2O [25] and EgoPAT3D [27], which contain egocentric RGB-D raw recordings for annotation purpose.

H2O [25] dataset is initially collected for 3D hand pose and interaction recognition using RGB-D data from both egocentric and multiple third-person views. We first use the precisely annotated 3D hand poses to compute the 3D centroids as the ground truth of the 3D hand trajectory, named **H2O-PT**, which is guaranteed to be of high-quality in [25] by multi-view verification.

EgoPAT3D [27] dataset is much larger than H2O. It is initially collected for predicting the 3D action targets from egocentric 3D videos. However, it does not provide either the 3D hand trajectory or the 3D hand poses. Thus, similar to [27], we develop an annotation workflow as shown in Fig. 4. More details about the annotation workflow are in the supplement. Eventually, we obtained sufficiently large

Table 1: **ADE and FDE results on H2O-PT dataset.** All models are built with ResNet-18 backbone. Best and secondary results are viewed in bold **black** and **blue** colors, respectively.

Model	ADE (\downarrow)			FDE (\downarrow)		
	3D _(3D)	2D _(3D)	2D _(2D)	3D _(3D)	2D _(3D)	2D _(2D)
DKF [24]	0.159	0.186	0.211	0.137	0.163	0.185
RVAE [26]	0.046	0.055	0.056	0.067	0.081	0.037
DSAE [48]	0.051	0.060	0.059	0.057	0.067	0.076
STORN [4]	0.043	0.053	0.053	0.094	0.141	0.076
VRNN [6]	0.041	0.050	0.050	0.051	0.081	0.068
SRNN [13]	0.040	0.048	0.049	0.036	0.061	0.044
EgoPAT3D* [27]	0.039	0.046	0.048	0.034	0.064	0.060
AGF* [50]	0.039	0.046	0.081	0.069	0.065	0.146
OCT* [34]	0.252	0.311	0.387	0.278	0.471	0.381
ProTran* [41]	0.066	0.088	0.109	0.099	0.168	0.123
USST	0.031	0.037	0.040	0.052	0.043	0.043

collections of 3D hand trajectories for training and evaluation, named **EgoPAT3D-DT**. To verify the reliability of the annotation workflow, we also apply it to H2O, resulting in a **H2O-DT** dataset.

Dataset Split The H2O(-PT/DT) dataset consists of 184 untrimmed long videos. We temporally sample the videos into multiple 64-frame clips with a step-size of 15 frames, resulting in 8203, 1735, and 3715 samples in training, validation, and testing splits, respectively. The EgoPAT3D-DT consists of 14 scenes and we split it into 11 seen scenes containing 8807 samples and 3 unseen testing scenes containing 2334 samples. The unseen scenes are not used in training, and the seen scenes are split into 6356, 846, and 1605 samples for training, validation, and seen testing.

Evaluation Setting We use the 3D Average Displacement Error (ADE) and Final Displacement Error (FDE) in meters as the evaluation metrics. The 2D trajectory results are normalized with reference to the video frame size. For all metrics, a small value indicates better performance. Each model is trained with 3D and 2D trajectory targets individually and evaluated by 3D metrics (3D_(3D)) and 2D metrics (2D_(2D)), respectively. The 3D model is additionally evaluated by 2D metrics (2D_(3D)) by projecting the 3D trajectory outputs to the 2D image plane.

4.2. Implementation Detail

The proposed method is implemented by PyTorch. In pre-processing, RGB videos are down-scaled to 64×64 . The 3D global trajectory data are normalized and further centralized to the range $[-1,1]$. By default, we set the observation ratio to 60%, the feature dimensions of $\mathbf{o}^{(V)}$ and $\mathbf{o}^{(T)}$ to 256, and the dimension of \mathbf{z} to 16 for all methods. In training, we use Huber loss to compute the location error.

Table 2: **ADE results on EgoPAT3D-DT dataset.** All models are built with ResNet-18 backbone. Best and secondary results are viewed in bold **black** and **blue** colors, respectively.

Model	Seen Scenes (\downarrow)			Unseen Scenes (\downarrow)		
	3D _(3D)	2D _(3D)	2D _(2D)	3D _(3D)	2D _(3D)	2D _(2D)
DKF [24]	0.294	0.237	0.157	0.260	0.202	0.133
RVAE [26]	0.216	0.110	0.121	0.194	0.104	0.109
DSAE [48]	0.214	0.129	0.143	0.188	0.116	0.131
STORN [4]	0.194	0.092	0.083	0.161	0.084	0.070
VRNN [6]	0.194	0.092	0.083	0.164	0.086	0.070
SRNN [13]	0.192	0.088	0.079	0.166	0.081	0.067
EgoPAT3D* [27]	0.186	0.081	0.079	0.170	0.080	0.068
AGF* [50]	6.149	0.136	0.099	6.045	0.119	0.087
OCT* [34]	0.853	0.163	0.098	0.782	0.139	0.091
ProTran* [41]	0.314	0.179	0.135	0.240	0.154	0.107
USST	0.183	0.089	0.082	0.120	0.075	0.060

We adopt the Adam optimizer with base learning rate $1e-4$ and cosine warmup scheduler for 500 training epochs on EgoPAT3D and 350 epochs on H2O datasets, respectively. More implementation details are in the supplement.

4.3. Main Results

Table 1 and 2 show the comparison between our method and existing sequential prediction approaches on H2O-PT and EgoPAT3D-DT datasets, respectively. The methods in the first multi-row section are general RNN-based while those in the second multi-row section show recent Transformer-based models. We put the FDE results on EgoPAT3D-DT in the supplement. The tables show our method achieves the best ADE performance and comparable FDE results with AGF and SRNN on H2O-PT, and significantly outperforms AGF, OCT, and ProTran on EgoPAT3D-DT. The competitive performance of SRNN is because of its both forward and backward passes over time such that all future positional encodings are utilized for forecasting. We notice that AGF, OCT, and ProTran do not work well on EgoPAT3D-DT, potentially due to the KL divergence vanish issue. The higher performance on the unseen split than the seen split can be attributed to the less distribution shift between unseen test trajectories and the training trajectories.

4.4. Model Analysis

Ablation Study To validate the effectiveness of the proposed modules and loss functions, we report the results of the ablation study in Table 3 on the EgoPAT3D-DT dataset.

We first compare the ProTran with the proposed state space transformer (SST), which is a vanilla version of USST

*We adapted the task-specific outputs of EgoPAT3D, AGF, OCT, and ProTran to fulfill the 3D hand trajectory forecasting task in this paper.

Table 3: **Ablation Study.** All models are trained with 3D targets and tested with both 3D and 2D ADE.

Variants	Seen (\downarrow)		Unseen (\downarrow)	
	3D	2D	3D	2D
ProTran (svi) [41]	0.314	0.179	0.240	0.154
ProTran (det) [41]	0.201	0.104	0.195	0.106
SST (ours)	0.190	0.088	0.174	0.084
USST w/o. $\mathcal{L}_{\text{haul}}$	0.292	0.237	0.267	0.204
USST w/o. $\mathcal{T}_{1:C}$	0.244	0.176	0.267	0.208
USST w/o. \mathbf{p}_{t-1}	0.196	0.090	0.169	0.098
USST w/o. $\mathcal{L}_{\text{velo}}$	0.189	0.091	0.168	0.099
USST w/o. w_t	0.183	0.090	0.130	0.077
USST (full model)	0.183	0.089	0.120	0.075

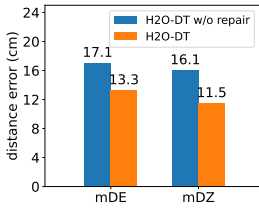


Figure 5: **Effect of depth repair.** mDE and mDZ are the mean error of 3D displacement and depth, respectively.

Metrics	SRNN		USST	
	3D _(3D)	2D _(2D)	3D _(3D)	2D _(2D)
H2O-PT	.087 / .076	.033 / .041	.040 / .049	.031 / .040
H2O-DT	.124 / .045	.052 / .041	.036 / .044	.052 / .043

Table 4: **Annotation Reliability.** ADE results (3D_(3D) / 2D_(2D)) are from testing on the same H2O-PT test set.

without uncertainty modeling, velocity constraint, and VPT. For a fair comparison, we implement a deterministic (det) version of ProTran in addition to the original method that uses stochastic variational inference (svi). Table 3 shows a clear advantage of our method over ProTran, which demonstrates the superiority of our SSM for state transition.

Next, in Table 3, we individually remove the new components and compare them with the full model of USST, including 1) the uncertainty loss function $\mathcal{L}_{\text{haul}}$, 2) the trajectory context $\mathcal{T}_{1:C}$, 3) the predictive link in $p_\phi(\mathbf{p}_t | \mathbf{z}_t, \mathbf{p}_{t-1})$, 4) the velocity constraint $\mathcal{L}_{\text{velo}}$, and 5) the depth robust weight w_t . It shows that uncertainty modeling is critical to guarantee reasonable forecasting results. Without historical trajectory $\mathcal{T}_{1:C}$, as expected, the performance degradation is significant. The predictive link from \mathbf{p}_{t-1} to \mathbf{p}_t is also important for the forecasting problem, which is consistent with the recent finding in [17]. It is noticeable that the velocity constraint shows a larger performance gain (4.8cm of 3D trajectory) on the unseen test set than on the seen data (0.6cm of 3D trajectory), revealing the importance of the physical rule for generalizable trajectory prediction. Lastly, the depth robust weight w_t (Eq. (10)) also boosts the performance of unseen data, showing the importance of modeling the depth noise from annotations.

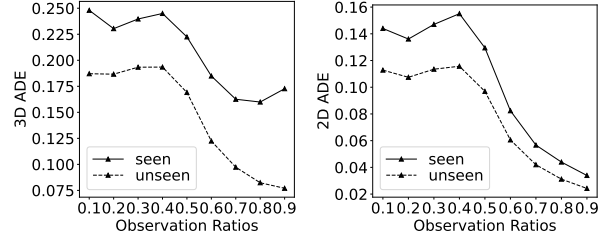


Figure 6: **Arbitrary Observation Ratios.** We report the results of 3D ADE (left) and 2D ADE (right) on EgoPAT3D-DT dataset.

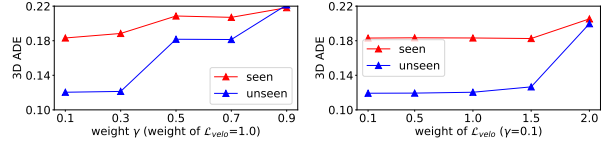


Figure 7: **Impact of loss weights.** Left: set the weight of $\mathcal{L}_{\text{velo}}$ to 1.0 and tune γ . Right: set γ to 0.1 and tune the weight of $\mathcal{L}_{\text{velo}}$.

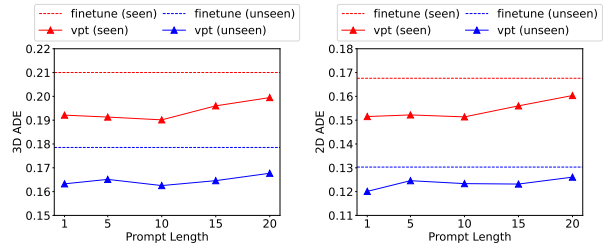


Figure 8: **Impact of Prompt Length.** We report the results of 3D (left) and 2D (right) ADE on EgoPAT3D-DT. The “finetune (unseen)” means finetune model on seen but test on unseen scenes.

Annotation Reliability By using the accurate H2O-PT as a reference, in Fig. 5, we show the effect of repairing depth of the 3D trajectory annotations from raw RGB-D data. We see a clear improvement in mDE and mDZ measurements. In Table 4, we further show the performance impact of annotation quality on SRNN and our USST models. It shows that our USST achieves more consistent ADE and FDE results than SRNN over the H2O-PT and H2O-DT. For reference, in the supplement, we additionally report the full results and analysis on H2O-DT and H2O-DT w/o repair.

Forecast at Any Time To simulate the real-world practice that forecasting trajectory at an arbitrary time, we take the advantage of the Transformer attention mask to fulfill random observation ratios ranging from 10% to 90%. The results are summarized in Fig. 6. It shows that with more percentage of information observed, both the 2D and 3D forecasting error are reduced as expected. It is interesting to see the slight increase of 3D ADE for the seen test data when using more observations. It could be caused by more inaccurate trajectory depth values at the end of trajectories.

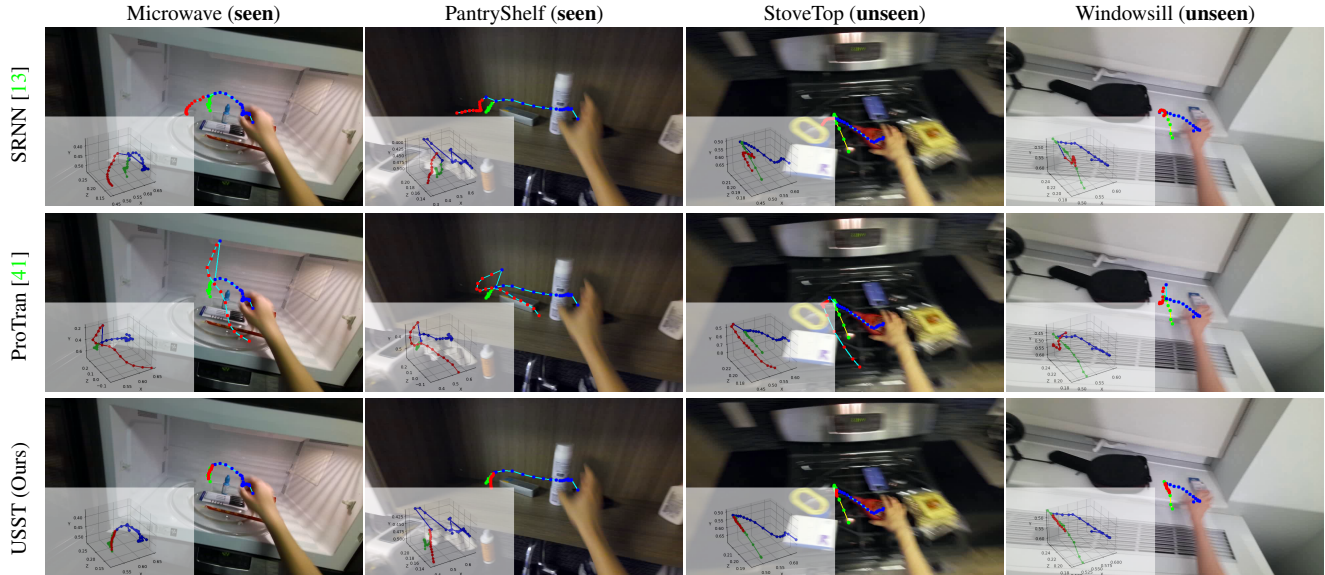


Figure 9: **Visualization on EgoPAT3D.** For each example (in a column), we show the global 3D trajectory and its 2D projection on the first frame. The blue, green, and red trajectory points represent the past observed, future ground truth, and future predictions, respectively.

Table 5: **Impact of 3D coordinate systems.** “Local” and “global” mean using 3D camera and world coordinates, respectively.

3D Target	Backbone	Seen (\downarrow)		Unseen (\downarrow)	
		3D	2D	3D	2D
Local	R18	0.202	0.083	0.174	0.062
Global	R18	0.183	0.089	0.120	0.075
Local	ViT	0.183	0.081	0.133	0.067
Global	ViT	0.182	0.087	0.119	0.075

Loss Weights Fig. 7 shows the EgoPAT3D-DT results of tuning the weights in Eq. (11), where the best performance is achieved when γ is set to 0.1 and the weight of $\mathcal{L}_{\text{velo}}$ is set to 1.0, respectively. We apply them to H2O-PT by default.

Prompt Length of VPT As indicated in VPT literature [21], the length of the visual prompt in ViT models needs to be carefully tuned for downstream tasks. In experiments, based on the SST model, we select the prompt length from $\{1, 5, 10, 15, 20\}$ and compare their performance with the baseline that fine-tunes the entire ViT backbone. Results are reported in Fig. 8. It shows that VPT could steadily achieve lower 2D and 3D ADE than fine-tuning, and the best performance is achieved when the prompt length is 10.

Local vs Global 3D Trajectory We note that the ambiguity of learning the appearance-location mapping exists when using local 3D targets. To justify the choice of global 3D trajectory targets, in Table 5, we compare the 3D and 2D ADE results using both ResNet-18 and ViT as visual backbones. It clearly shows that for 3D trajectory prediction, a

global 3D coordinate system is a better choice, while for 2D trajectory evaluation, the local 3D target is better. These observations are expected as in the local 3D coordinate system, the projected 2D pixel locations of moving hands tend to be in the visual center due to the egocentric view so that the model training is dominated by the 2D hand locations.

Qualitative Results As shown in Fig. 9, the proposed USST model is compared with the Transformer-based approach ProTran and the most competitive method SRNN. It clearly shows that our trajectory forecasting is much better than the three compared methods. Please refer to our supplement for more visualizations.

Limitations & Discussions The dataset annotation is limited in scenarios when the RGB-D sensors or camera poses are not available. The model is limited in the recursive way of state transition, which is not hardware friendly for parallel inference. Besides, in the future, other tasks like the 3D hand pose and interaction recognition can be jointly studied for a fine-grained egocentric understanding.

5. Conclusion

In this paper, we propose to forecast human hand trajectory in 3D physical space from egocentric videos. For this goal, we first develop a pipeline to automate the 3D trajectory annotation. Then, we propose a novel uncertainty-aware state space transformer (USST) model to fulfill the task. Empirically, with the aleatoric uncertainty modeling, velocity constraint, and visual prompt tuning, our model achieves the best performance on both H2O and EgoPAT3D datasets and good generalization to the unseen scenes.

References

- [1] 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. [2](#)
- [2] Wentao Bao, Bin Xu, and Zhenzhong Chen. Monofenet: Monocular 3d object detection with feature enhancement networks. *IEEE TIP*, 29:2753–2765, 2019. [2](#)
- [3] Wentao Bao, Qi Yu, and Yu Kong. Uncertainty-based traffic accident anticipation with spatio-temporal relational learning. In *ACM MM*, 2020. [2, 4](#)
- [4] Justin Bayer and Christian Osendorfer. Learning stochastic recurrent networks. In *NeurIPS Workshop*, 2014. [2, 6](#)
- [5] Huikun Bi, Ruisi Zhang, Tianlu Mao, Zhigang Deng, and Zhaoqi Wang. How can i see my future? fvtraj: Using first-person view for pedestrian trajectory prediction. In *ECCV*, 2020. [1, 2](#)
- [6] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratharth Goel, Aaron C Courville, and Yoshua Bengio. A recurrent latent variable model for sequential data. In *NeurIPS*, 2015. [2, 4, 6](#)
- [7] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *ECCV*, 2018. [1, 2](#)
- [8] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. The epic-kitchens dataset: Collection, challenges and baselines. *IEEE TPAMI*, 43(11):4125–4141, 2020. [1, 2](#)
- [9] Nachiket Deo and Mohan M Trivedi. Convolutional social pooling for vehicle trajectory prediction. In *CVPR*, 2018. [2](#)
- [10] Christian Diller, Thomas Funkhouser, and Angela Dai. Forecasting characteristic 3d poses of human actions. In *CVPR*, 2022. [1](#)
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2020. [4, 5](#)
- [12] Alireza Fathi, Ali Farhadi, and James M Rehg. Understanding egocentric activities. In *ICCV*, 2011. [2](#)
- [13] Marco Fraccaro, Søren Kaae Sønderby, Ulrich Paquet, and Ole Winther. Sequential neural models with stochastic layers. In *NeurIPS*, 2016. [2, 4, 6, 8](#)
- [14] Hao Fu, Chunyuan Li, Xiaodong Liu, Jianfeng Gao, Asli Celikyilmaz, and Lawrence Carin. Cyclical annealing schedule: A simple approach to mitigating KL vanishing. In *NACCL*, 2019. [3, 4](#)
- [15] Nisal Menuka Gamage, Deepana Ishtaweera, Martin Weigel, and Anusha Withana. So predictable! continuous 3d hand trajectory prediction in virtual reality. In *UIST*, 2021. [1](#)
- [16] Rohit Girdhar and Kristen Grauman. Anticipative video transformer. In *ICCV*, 2021. [1](#)
- [17] Laurent Girin, Simon Leglaive, Xiaoyu Bie, Julien Diard, Thomas Hueber, and Xavier Alameda-Pineda. Dynamical variational autoencoders: A comprehensive review. *Foundations and Trends in Machine Learning*, 15(1-2):1–175, 2021. [2, 7](#)
- [18] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *CVPR*, 2022. [1, 2](#)
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. [4, 5](#)
- [20] Yihui He, Chenchen Zhu, Jianren Wang, Marios Savvides, and Xiangyu Zhang. Bounding box regression with uncertainty for accurate object detection. In *CVPR*, 2019. [3, 4](#)
- [21] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *ECCV*, 2022. [5, 8](#)
- [22] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *NeurIPS*, 2017. [3, 4](#)
- [23] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *CVPR*, 2018. [4](#)
- [24] Rahul G Krishnan, Uri Shalit, and David Sontag. Deep kalman filters. *arXiv preprint arXiv:1511.05121*, 2015. [6](#)
- [25] Taemin Kwon, Bugra Tekin, Jan Stühmer, Federica Bogo, and Marc Pollefeys. H2O: Two hands manipulating objects for first person interaction recognition. In *ICCV*, 2021. [2, 5](#)
- [26] Simon Leglaive, Xavier Alameda-Pineda, Laurent Girin, and Radu Horaud. A recurrent variational autoencoder for speech enhancement. In *ICASSP*, 2020. [6](#)
- [27] Yiming Li, Ziang Cao, Andrew Liang, Benjamin Liang, Luoyao Chen, Hang Zhao, and Chen Feng. Egocentric prediction of action target in 3d. In *CVPR*, 2022. [1, 2, 4, 5, 6](#)
- [28] Yin Li, Miao Liu, and James M. Rehg. In the eye of beholder: Joint learning of gaze and actions in first person video. In *ECCV*, 2018. [1](#)
- [29] Yanghao Li, Tushar Nagarajan, Bo Xiong, and Kristen Grauman. Ego-exo: Transferring visual representations from third-person to first-person videos. In *CVPR*, 2021. [1, 2](#)
- [30] Zhengqi Li, Tali Dekel, Forrester Cole, Richard Tucker, Noah Snavely, Ce Liu, and William T Freeman. Learning the depths of moving people by watching frozen people. In *CVPR*, 2019. [2](#)
- [31] Miao Liu, Lingni Ma, Kiran Somasundaram, Yin Li, Kristen Grauman, James M Rehg, and Chao Li. Egocentric activity recognition and localization on a 3d map. In *ECCV*, 2022. [1](#)
- [32] Miao Liu, Siyu Tang, Yin Li, and James M Rehg. Forecasting human-object interaction: joint prediction of motor attention and actions in first person video. In *ECCV*, 2020. [1](#)
- [33] Miao Liu, Siyu Tang, Yin Li, and James M Rehg. Forecasting human-object interaction: joint prediction of motor attention and actions in first person video. In *ECCV*, 2020. [2](#)
- [34] Shaowei Liu, Subarna Tripathi, Somdeb Majumdar, and Xiaolong Wang. Joint hand motion and interaction hotspots prediction from egocentric videos. In *CVPR*, 2022. [1, 2, 4, 6](#)

- [35] Srikanth Malla, Isht Dwivedi, Behzad Dariush, and Chiho Choi. Nemo: Future object localization using noisy ego priors. In *ITSC*, 2022. [2](#)
- [36] Tushar Nagarajan, Yanghao Li, Christoph Feichtenhofer, and Kristen Grauman. Ego-topo: Environment affordances from egocentric video. In *CVPR*, 2020. [1](#)
- [37] Adrián Núñez-Marcos, Gorka Azkune, and Ignacio Arganda-Carreras. Egocentric vision-based action recognition: a survey. *Neurocomputing*, 472:175–197, 2022. [2](#)
- [38] Paschalis Panteleris, Iason Oikonomidis, and Antonis Argyros. Using a single rgb frame for real time 3d hand pose estimation in the wild. In *WACV*, 2018. [2](#)
- [39] Jianing Qiu, Lipeng Chen, Xiao Gu, Frank P-W Lo, Ya-Yen Tsai, Jiankai Sun, Jiaqi Liu, and Benny Lo. Egocentric human trajectory forecasting with a wearable camera and multi-modal fusion. *IEEE Robotics and Automation Letters*, 7(4):8799–8806, 2022. [2](#)
- [40] Syama Sundar Rangapuram, Matthias W Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, and Tim Januschowski. Deep state space models for time series forecasting. In *NeurIPS*, 2018. [2](#), [3](#)
- [41] Binh Tang and David S Matteson. Probabilistic transformer for time series analysis. In *NeurIPS*, 2021. [3](#), [4](#), [6](#), [7](#), [8](#)
- [42] Bugra Tekin, Federica Bogo, and Marc Pollefeys. H+o: Unified egocentric recognition of 3d hand-object poses and interactions. In *CVPR*, 2019. [2](#)
- [43] Luca Anthony Thiede and Pratik Prabhanjan Brahma. Analyzing the variety loss in the context of probabilistic trajectory prediction. In *CVPR*, 2019. [2](#)
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. [3](#)
- [45] Jiashun Wang, Huazhe Xu, Medhini Narasimhan, and Xiaolong Wang. Multi-person 3d motion prediction with multi-range transformers. In *NeurIPS*, 2021. [1](#)
- [46] Xiaohan Wang, Yu Wu, Linchao Zhu, and Yi Yang. Symbiotic attention with privileged information for egocentric action recognition. In *AAAI*, 2020. [1](#)
- [47] Jürgen Wiest, Matthias Höffken, Ulrich Kreßel, and Klaus Dietmayer. Probabilistic trajectory prediction with gaussian mixture models. In *IVS*, 2012. [2](#)
- [48] Li Yingzhen and Stephan Mandt. Disentangled sequential autoencoder. In *ICML*, 2018. [6](#)
- [49] Ye Yuan and Kris Kitani. Ego-pose estimation and forecasting as real-time PD control. In *ICCV*, 2019. [1](#)
- [50] Ye Yuan, Xinshuo Weng, Yanglan Ou, and Kris M Kitani. Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting. In *ICCV*, 2021. [3](#), [6](#)
- [51] Pu Zhang, Wanli Ouyang, Pengfei Zhang, Jianru Xue, and Nanning Zheng. Sr-lstm: State refinement for lstm towards pedestrian trajectory prediction. In *CVPR*, 2019. [2](#)
- [52] Zhoutong Zhang, Forrester Cole, Richard Tucker, William T Freeman, and Tali Dekel. Consistent depth of moving objects in video. *ACM TOG*, 40(4):1–12, 2021. [2](#)