

# COPILLOT: Human-Environment Collision Prediction and Localization from Egocentric Videos

Boxiao Pan<sup>1</sup> Bokui Shen<sup>1\*</sup> Davis Rempe<sup>1,2\*</sup> Despoina Paschalidou<sup>1</sup>  
 Kaichun Mo<sup>1,2</sup> Yanchao Yang<sup>1,3</sup> Leonidas J. Guibas<sup>1</sup>  
<sup>1</sup>Stanford University <sup>2</sup>NVIDIA Research <sup>3</sup>The University of Hong Kong

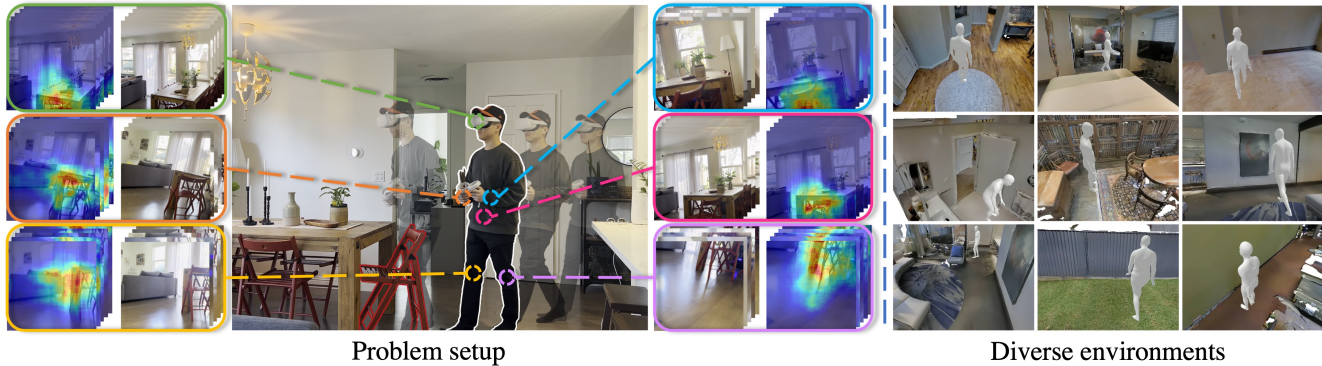


Figure 1: **Problem formulation.** (Left) Egocentric videos from cameras mounted on the body are used to perform three tasks: binary collision prediction, classifying the colliding joints, and localizing collision regions in the environment. In this example, our model predicts and localizes a collision between the left knee and the chair. (Right) Our model is trained on a newly proposed large-scale synthetic dataset and generalizes to diverse synthetic and real-world environments.

## Abstract

The ability to forecast human-environment collisions from egocentric observations is vital to enable collision avoidance in applications such as VR, AR, and wearable assistive robotics. In this work, we introduce the challenging problem of predicting collisions in diverse environments from multi-view egocentric videos captured from body-mounted cameras. Solving this problem requires a generalizable perception system that can classify which human body joints will collide and estimate a collision region heatmap to localize collisions in the environment. To achieve this, we propose a transformer-based model called COPILLOT to perform collision prediction and localization simultaneously, which accumulates information across multi-view inputs through a novel 4D space-time-viewpoint attention mechanism. To train our model and enable future research on this task, we develop a synthetic data generation framework that produces egocentric videos of virtual humans moving and colliding within diverse 3D environments. This framework is then used to establish a large-scale dataset consisting of 8.6M egocentric RGBD frames. Extensive experiments show that COPILLOT generalizes to unseen synthetic as well as real-world scenes. We further

demonstrate COPILLOT outputs are useful for downstream collision avoidance through simple closed-loop control. Please visit our project webpage at <https://sites.google.com/stanford.edu/copilot>.

## 1. Introduction

Forecasting potential collisions between a human and their environment from *egocentric* observations is crucial to applications such as virtual reality (VR), augmented reality (AR), and wearable assistive robotics. As users are immersed in VR / AR, the effective and non-intrusive prevention of collisions with real-world surroundings is critical. On the other hand, wearable exoskeletons [52, 10] provide physical assistance to users with mobility impairments and therefore must control user movements to avoid dangerous falls by estimating *when* and *where* a collision may occur.

To this end, a multitude of works have explored ways to predict and localize future collisions. For example, existing approaches for VR leverage redirected walking techniques [42] to guide users in the virtual world such that they avoid collisions with real-world objects [23, 22, 48, 15, 57]. However, these approaches rely on pre-scanned scenes and external tracking systems [23, 48, 15, 57] or resort to on-line SLAM [22, 12], making them cumbersome and tedious to use. For assistive exoskeletons, several methods have

\* Equal contribution

been proposed to classify environment types [28, 27], recognize obstacle parameters [31, 25], predict collision risks [26, 3], and predict future trajectories [5, 46]. While these works are useful in constrained settings, they impose several constraints to simplify the problem, such as assuming pre-defined environments [3, 25, 27] or performing only environment classification rather than explicit collision forecasting [40, 28, 41, 11].

In this work, we address the shortcomings of prior approaches to collision forecasting by introducing a generalizable model that provides detailed perception outputs useful for downstream collision avoidance. We first formalize the problem of egocentric collision prediction and localization from single or multi-view videos captured by cameras mounted on a person. To tackle this problem, we design a multi-task architecture that scales to diverse and complex environments and provides rich collision information. We train our model using high-quality simulation and curate data from a large variety of synthetic environments.

We pose collision forecasting as a problem of *classifying* which body joints will collide in the near future and *localizing* regions in the input videos where these collisions are likely to occur. The input is a set of multi-view RGB egocentric videos (see Fig. 1) *without* camera poses, which makes our problem setup general but also more challenging. To successfully predict when and what part of the human body will collide with the environment, we must accumulate information about the human’s movement, intent, and their surroundings across the multi-view videos.

To address these challenges, we propose COPILOT, a **C**ollision **P**red**I**ction and **L**ocalization **T**ransformer that uses a novel 4D attention scheme across space, time, and viewpoint to leverage information from multi-view video inputs. By alternating between cross-viewpoint and cross-time attention, the proposed 4D attention fuses information on both scene geometry and human motion. Notably, COPILOT classifies collisions and estimates collision regions simultaneously in a multi-task fashion, thereby improving performance on collision prediction and providing actionable information for collision avoidance as well.

To train and evaluate our model, we develop a data framework and curate a high-quality synthetic dataset containing  $\sim 8.6$ M egocentric RGBD frames with automatically-annotated collision labels and heatmaps. Data is generated in simulation [45, 47] using real-world scene scans [8, 60] and a human motion generation model pre-trained with motion capture data [43]. Our data framework and dataset feature large scene diversity, realistic human motion, and accurate collision checking. Experiments show that COPILOT achieves over 84% collision prediction accuracy on unseen synthetic scenes after training on our large-scale dataset, and qualitatively generalizes to diverse real-world environments. To mimic the exoskeleton

application, we combine our model with a closed-loop controller in simulation to demonstrate that we can adjust human motion for improved collision avoidance; this avoids 35% and 29% of collision cases on training and unseen synthetic scenes, respectively.

In summary, we make the following **contributions**: (1) we introduce the challenging task of collision prediction and localization from unposed egocentric videos. (2) We propose a multi-view video transformer-based model that employs a novel 4D space-time-viewpoint attention mechanism to solve this task effectively. (3) We develop a synthetic data framework to generate a large-scale dataset of egocentric RGBD frames with automatically-annotated collision labels and heatmaps. Our model trained on this dataset is shown to generalize to diverse synthetic and real-world environments through extensive evaluations. The dataset and code will be made available upon acceptance.

## 2. Related Work

**Collision perception for VR.** To avoid potential user collisions with the real-world environment, existing works combine Redirected Walking (RDW) [42] with techniques such as Artificial Potential Fields [48, 15] and Reinforcement Learning [29, 9]. On a high level, RDW maps the obstacles in the physical world to some form in the virtual world to imperceptibly steer the user along a collision-free path. These techniques require known scene geometries and user locations in the physical world, which normally come from scene scanning and external tracking systems. Some works try to relax these assumptions by fusing depth observations online [22, 12], but they require additional hardware and tedious post-processing. As our approach directly predicts collisions from egocentric video streams, it does not require cumbersome external hardware setups.

**Collision perception for exoskeletons.** Assistive wearable exoskeletons provide physical assistance to modify a user’s actions and avoid collisions with the environment. To achieve this, environment perception is necessary. Prior works classify environments into one of several pre-defined categories [28, 41, 27], recognize obstacle parameters [31, 25], or predict short-term future user motions [3, 5, 46]. These estimations are fed to hand-designed algorithms to guide the behavior of the exoskeleton controller. In contrast, our approach explicitly forecasts collisions, without relying on any post-processing. Moreover, these works operate in simple environments with few obstacles and only consider the lower body limbs, unlike COPILOT that generalizes to diverse environments and predicts full-body collisions.

**Perception from egocentric videos.** While egocentric vision spans a diverse range of topics, the most relevant to our work are body pose estimation, human-object interaction prediction, and collision prediction. Estimating and fore-

casting full-body poses from egocentric videos is an under-constrained problem since multiple human poses can plausibly explain the same egocentric observations. Prior works use a downward-facing fish-eye camera to better capture body parts [49, 53], regularize estimations with dynamics and kinematics models [61, 36], or leverage multi-person interaction cues in the egocentric view [39]. In this paper, we do not focus on this task but instead seek to exploit cues from the input egocentric videos to reason about potential collisions with the environment in the near future. Other works infer human-object interactions from egocentric videos, such as future 2D hand trajectories and interaction hotspots [32, 33]. The closest to ours is [38], which predicts the time-to-environment collision for a suitcase-shaped robot, given only a single-view RGB video. In contrast, we predict fine-grained per-joint collisions and collision heatmaps for the more complex human morphology.

**General video understanding.** Spatio-temporal reasoning is one of the main topics for video understanding. Earlier approaches perform 3D convolutions to allow hierarchical spatio-temporal understanding [50, 7, 20, 56, 19, 58]. Recently, several works [2, 18, 34, 6, 37] have adapted powerful transformer architectures [51] to the video domain. Similar to ViT [16], TimeSformer [6] divides video frames into patches and proposes several strategies to perform spatio-temporal self-attention. This architecture achieves strong performance on capturing spatio-temporal patterns in video data, but only handles a single video stream. Our backbone extends TimeSformer to use a novel 4D attention operating over the space-time-viewpoint domain.

### 3. Problem Formulation and Data Generation

#### 3.1. Collision Prediction and Localization

To provide useful information for downstream collision avoidance, we need to know *if* there will be a collision between the person and the environment as well as *where* this collision will occur, both on the human body and in the scene. Given multi-view egocentric videos, we formulate this problem as (1) classifying whether a collision will happen in the near future, (2) identifying the body joints involved in a potential collision, and (3) localizing the collision region in the scene. Unlike prior works that investigate collisions only for lower limbs [28, 25, 46, 52, 27, 3, 5], we expand our scope to full-body collision prediction that utilizes multi-view inputs from cameras mounted on several body joints. Moreover, as camera extrinsics and intrinsics can be hard to acquire for in-the-wild scenes, we rely solely on visual inputs. These unposed multi-view input videos have the potential to capture the full-body motion and environment context better than a single view.

Given a set of multi-view, time-synchronized egocentric RGB videos  $X = \{X_1, X_2, \dots, X_V\}$  captured from

$V$  viewpoints where  $X_v = \{x_1^v, x_2^v, \dots, x_T^v\}$  consists of  $T$  frames captured from the  $v$ -th camera, our goal is to predict:

1. An *overall* binary prediction  $y^{\text{col}}$  indicating whether any collision will happen in the next  $H$  time steps. This determines *whether* to intervene to avoid a collision;
2. A *per-joint* binary prediction  $y_j^{\text{joint}}$  indicating whether the  $j$ -th joint will be involved in a collision. Note that typically, the number of joints  $J$  is different from the number of views  $V$ , because attaching cameras to all body joints can be impractical. This indicates *where* to intervene;
3. A set of per-frame collision region heatmaps  $y^{\text{map}} = \{P_1, P_2, \dots, P_V\}$ , where  $P_v = \{p_1^v, p_2^v, \dots, p_T^v\}$  is the heatmap sequence corresponding to  $X_v$ . Note that the sum of all values in one heatmap is one, namely  $\sum_{ij} p_{ij}^v = 1$ , so that  $p_i^v$  is a probability distribution over all spatial locations in that frame. This represents the likelihood that a collision will occur at each location and indicates *how* to intervene.

Although in the above formulation we only consider RGB inputs, leveraging depth data is also possible. However, we choose RGB images as our primary setting because they are easier to acquire than depth images, which rely on cumbersome depth cameras that are not always available.

#### 3.2. Large-Scale Data Generation

Due to the anomalous nature of collision events, it is unclear how to collect real-world data for our task, let alone fine-grained annotations for collision regions and colliding joints. Hence, we build a framework that leverages a simulated environment to generate a large-scale high-quality synthetic dataset containing  $\sim 8.6\text{M}$  egocentric RGBD frames with automatically-annotated collision labels and heatmaps. We describe the key steps of this framework next and leave the details to the supplementary material.

Our data framework places virtual humans in synthetic scenes and renders egocentric videos as they perform realistic motions. To get a diverse set of environments, we pick 100 scenes from the Matterport3D [8] and Gibson [60] datasets, where each scene contains several rooms. We leverage HuMoR [43] to randomly generate sequences of realistic walking motions in each scene using the SMPL human body model [35]. During sequence rollout, we perform collision checking between the human and scene meshes to get the *overall* collision label  $y^{\text{col}}$ . We then assign each vertex where a collision occurred to the closest one of 10 human body joints (head, torso, elbows, hands, legs, and feet), which results in a set of *per-joint* collision labels  $y^{\text{joint}}$ . Collision checking is efficiently implemented using PyBullet [14] by segmenting the body into convex parts; when a collision is detected, the sequence is terminated.

For each generated motion sequence, observations are rendered from egocentric viewpoints with the AI-habitat

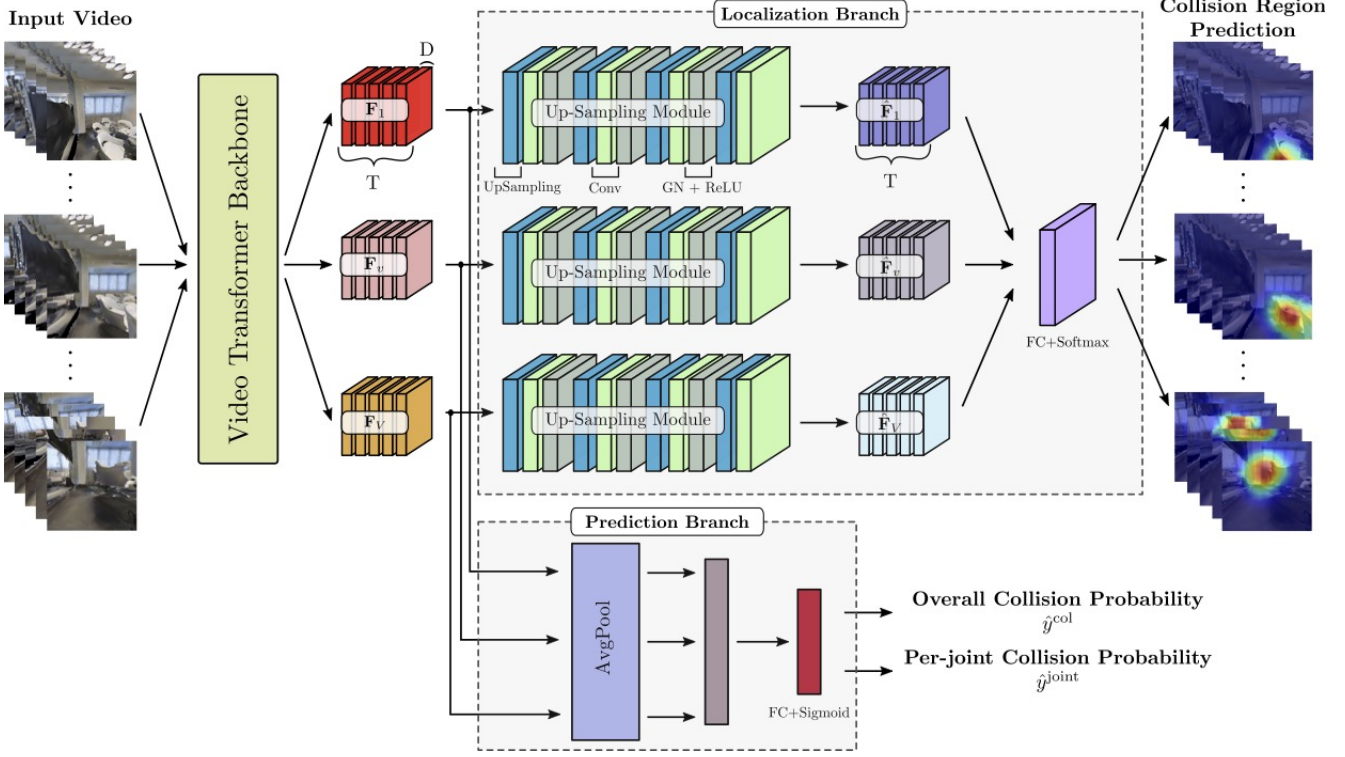


Figure 2: **COPILOT architecture**. Videos taken from different viewpoints are fed to a backbone transformer model that uses 4D space-time-viewpoint attention to output a set of feature maps. The localization branch gradually up-samples the features to the original image resolution via the up-sampling module, which are then used to estimate the collision region heatmaps. The prediction branch pools the features globally to predict the overall and joint-level collisions.

simulator [47, 45] by placing cameras at 6 body joints (head, pelvis, wrists, and knees). Each sequence is split into subsequences of one second length that are used as input to our model; each subsequence is labeled as “colliding” if a collision is detected within one second in the future (*i.e.*  $H = 30$  at 30 fps). We empirically found (Sec. 5.4) that videos of a one-second duration typically have enough evidence to decide whether to act to avoid a potential collision. Finally, to generate the collision region heatmaps  $y^{\text{map}}$ , we project the vertices in the scene that the person collided with to each egocentric viewpoint, and set these pixels to 1 and the rest to 0. We apply a Gaussian kernel to smooth the heatmap before normalizing it to a proper distribution.

## 4. Method

Our goal is to develop a perception module that can jointly solve the collision prediction and collision region localization tasks. As these tasks are deeply connected and both require learning features of motion and environment geometry, we introduce COPILOT to tackle them jointly in a multi-task fashion (supported experimentally in Sec. 5.1.2). As shown in Fig. 2, the input egocentric videos are first fed into a transformer encoder, which performs

attention over space, time, and viewpoint across all video streams using our novel 4D attention mechanism. Next, the *localization branch* takes the output features and estimates collision region heatmaps for each viewpoint, while the *prediction branch* performs global pooling on the features and predicts the overall and per-joint collision probabilities.

### 4.1. Transformer Backbone with 4D Attention

Our backbone is based on TimeSformer [6], which we extend to process videos from several views jointly using cross-viewpoint attention. Following the original architecture [6], the model first splits each  $H \times W$  RGB frame into  $N$  non-overlapping spatial patches, each of size  $P \times P$ . Each patch is flattened into a vector  $\mathbf{z}_{(p,t,v)} \in \mathbb{R}^F$ , where  $p$ ,  $t$ , and  $v$  index space, time, and viewpoint, respectively, and  $F = 3 \cdot P \cdot P$ . Each vector is then mapped to a *patch embedding*  $\mathbf{e}_{(p,t,v)}^{(0)} \in \mathbb{R}^D$  as

$$\mathbf{e}_{(p,t,v)}^{(0)} = E\mathbf{z}_{(p,t,v)} + \text{emb}_{(p,t,v)}^{\text{pos}}, \quad (1)$$

where  $E \in \mathbb{R}^{D \times F}$  is a learnable matrix and  $\text{emb}_{(p,t,v)}^{\text{pos}} \in \mathbb{R}^D$  is a learnable positional embedding. These patch embeddings are the input to the transformer encoding blocks.

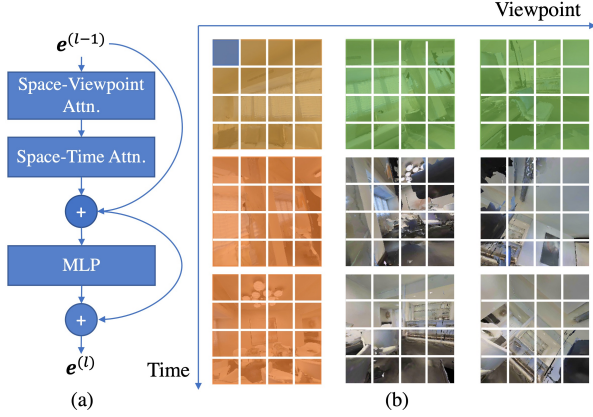


Figure 3: **4D STV attention.** (a) Attention block overview. (b) Illustration of the attention operation. Patches from three time steps and viewpoints are visualized. Green and orange denote patches that the blue patch attends to along the space-viewpoint and space-time dimensions, respectively. The current frame (upper left) has both colors.

In particular, for a single encoding block  $l \in [1, L]$  (see Fig. 3), the query vectors are computed as

$$\mathbf{q}_{(p,t,v)}^{(l,a)} = W_Q^{(l,a)} \text{LN} \left( \mathbf{e}_{(p,t,v)}^{(l-1)} \right), \quad (2)$$

where  $\text{LN}(\cdot)$  denotes LayerNorm [4] and  $a \in [1, \mathcal{A}]$  indexes the attention heads. Similarly, we compute the keys  $\mathbf{k}_{(p,t,v)}^{(l,a)}$  and values  $\mathbf{v}_{(p,t,v)}^{(l,a)}$ .

TimeSformer [6] proposes several attention mechanisms to perform space-time self-attention on a single video stream. One way to adapt their model to our multi-view setting is to perform space-time attention for each stream separately and then fuse the resulting features with an MLP (*i.e.* late feature fusion). However, this strategy does not fully exploit the overlapping multi-view information relevant to scene geometry and human motion. We thus propose a 4D space-time-viewpoint (STV) attention operation, which is illustrated in Fig. 3. For each patch, the model first attends to all patches across different viewpoints at the same time step (*i.e.* space-viewpoint attention), then attends to all patches within the same viewpoint (*i.e.* space-time attention). Intuitively, space-viewpoint attention gathers environment information from all viewpoints, while space-time attention captures dynamics from a single viewpoint.

Formally, we first compute self-attention weights  $\alpha_{\text{SV}} \in \mathbb{R}^{N \cdot V}$  and  $\alpha_{\text{ST}} \in \mathbb{R}^{N \cdot T}$  for each attention type:

$$\alpha_{\text{SV}(p,t,v)}^{(l,a)} = \text{SM} \left( \frac{\mathbf{q}_{(p,t,v)}^{(l,a) \top}}{\sqrt{D_h}} \cdot \left\{ \mathbf{k}_{(p',t,v')}^{(l,a)} \right\}_{\substack{p'=1,\dots,N \\ v'=1,\dots,V}} \right) \quad (3)$$

$$\alpha_{\text{ST}(p,t,v)}^{(l,a)} = \text{SM} \left( \frac{\mathbf{q}_{(p,t,v)}^{(l,a) \top}}{\sqrt{D_h}} \cdot \left\{ \mathbf{k}_{(p',t',v)}^{(l,a)} \right\}_{\substack{p'=1,\dots,N \\ t'=1,\dots,T}} \right) \quad (4)$$

where  $\text{SM}(\cdot)$  denotes the softmax function. The output feature encoding  $\mathbf{e}_{(p,t,v)}^{(l)}$  is computed by first summing the value vectors with the attention weights:

$$\mathbf{s}_{(p,t,v)}^{(l,a)} = \sum_{p'=1}^N \sum_{v'=1}^V \alpha_{\text{SV}(p,t,v),(p',v')}^{(l,a)} \mathbf{v}_{(p',v',t)}^{(l,a)} \quad (5)$$

$$+ \sum_{p'=1}^N \sum_{t'=1}^T \alpha_{\text{ST}(p,t,v),(p',t')}^{(l,a)} \mathbf{v}_{(p',v,t')}^{(l,a)}. \quad (6)$$

Then, the encodings from all heads are concatenated and mapped through an MLP with residual connections to obtain the final feature output of the block:

$$\mathbf{e}_{(p,t,v)}^{(l)} = W_O \left[ \mathbf{s}_{(p,t,v)}^{(l,1)} \dots \mathbf{s}_{(p,t,v)}^{(l,\mathcal{A})} \right] + \mathbf{e}_{(p,t,v)}^{(l-1)} \quad (7)$$

$$\mathbf{e}_{(p,t,v)}^{(l)} = \text{MLP} \left( \text{LN} \left( \mathbf{e}_{(p,t,v)}^{(l)} \right) \right) + \mathbf{e}_{(p,t,v)}^{(l)}. \quad (8)$$

## 4.2. Collision Region Localization

The output feature grid is then passed to the *localization branch*, which predicts the per-frame spatial heatmaps corresponding to the input videos. High values in the heatmap indicate areas in the scene where there is a high probability for a collision to occur in the near future. Inferring spatial collision heatmaps requires per-pixel predictions, so this branch of the model must up-sample the discretized feature maps back to the original image resolution. Inspired by the scheme introduced in FPN [30], we use a sequence of up-sampling modules, each doing  $2 \times$  nearest-neighbor up-sampling followed by a  $3 \times 3$  convolution, GroupNorm layer [59], and ReLU activation. Different from [30], the feature dimension is reduced by half in each spatial up-sampling to reduce the memory allocation. Finally, we use a one-layer MLP followed by a *softmax* to get the collision region heatmaps  $\hat{y}_v^{\text{map}} \in [0, 1]^{H \times W}$ , such that each heatmap is a distribution of collision probability:

$$\hat{y}_v^{\text{map}} = \text{SM} \left( \text{MLP} \left( \text{up} \left( \mathbf{e}_{(t,v)}^{(L)} \right) \right) \right), \quad (9)$$

where  $\mathbf{e}_{(t,v)}^{(L)}$  is the  $(t, v)$ -th slice of  $\mathbf{e}^{(L)}$ .

## 4.3. Overall and Per-Joint Collision Prediction

Likewise, we pass the output features from our transformer backbone to the *prediction branch* that predicts: (1) the *overall* binary collision probability, namely whether the person will collide within the next  $H$  timesteps, and (2) the *per-joint* binary collision probability indicating whether each joint will collide. Features are first pooled globally using average-pooling across all space-time-viewpoint. Then, this feature vector is fed to an MLP followed by a sigmoid activation to estimate the overall and per-joint collision predictions as follows:

$$\hat{y}^{\text{joint}}, \hat{y}^{\text{col}} = \text{Sig} \left( \text{MLP} \left( \text{pool} \left( \mathbf{e}^{(L)} \right) \right) \right) \in [0, 1]^{J+1}, \quad (10)$$

where the first  $J$  outputs from the MLP are used for joint collision prediction  $\hat{y}^{\text{joint}} \in [0, 1]^J$  while the last value is used for overall collision prediction  $\hat{y}^{\text{col}} \in [0, 1]$ .

#### 4.4. Training

COPILLOT is trained in a fully-supervised manner using the synthetic dataset described in Sec. 3.2. The overall loss is the sum of three terms

$$L = L_{\text{map}} + L_{\text{col}} + L_{\text{joint}}, \quad (11)$$

where  $L_{\text{map}} = D_{\text{KL}}(\hat{y}^{\text{map}} || y^{\text{map}})$  is the KL-divergence between the predicted and the ground-truth collision heatmaps,  $L_{\text{col}} = \text{BCE}(\hat{y}^{\text{col}}, y^{\text{col}})$  is the binary cross-entropy loss (BCE) between the predicted and the target overall collision label for the motion sequence, and  $L_{\text{joint}} = \sum_{j=1}^J \text{BCE}(\hat{y}_j^{\text{joint}}, y_j^{\text{joint}})$  is the sum of the binary cross-entropy losses between the target and the predicted per-joint collision probabilities.

**Implementation Details.** We use PyTorch Lightning [17] and train using SGD [44] with an initial learning rate of  $1e-4$ , which is decayed by 0.1 at the 25-th, 35-th, and 45-th epochs. All training is done on a single Nvidia RTX A5000 GPU, and the model requires training for approximately two days to converge. For all experiments, the model is given 1 sec of past observations (sampled at 5 Hz) and predicts collisions up to 1 sec into the future.

### 5. Experimental Results

We compare COPILLOT to baselines in both single-view and multi-view settings in Sec. 5.1.1. In Sec. 5.1.2, we investigate the impact of various components of COPILLOT on its performance. We then present qualitative evaluation results of collision region localization, both in simulation (Sec. 5.2) and the real world (Sec. 5.3). Additionally, we show that outputs from COPILLOT are useful by applying COPILLOT to a collision avoidance task. The supplementary material provides additional details of experiments and metrics, along with additional visualizations and analyses.

#### 5.1. Quantitative Evaluation

We evaluate COPILLOT and baselines on unseen motion sequences from the synthetic dataset described in Sec. 3.2. We consider two evaluation setups. In the first, we evaluate on training scenes but novel motion sequences, whereas in the second, both scenes and motion sequences are not seen during training. We refer to the former as *Unseen Motions* and to the latter as *Unseen Motions & Scenes*. To measure the quality of the overall collision prediction  $\hat{y}^{\text{col}}$ , we report the binary classification accuracy (*Col*). For the per-joint collision predictions  $\hat{y}^{\text{joint}}$ , we report F1-score ( $F1@N$ ,  $N=4,6,10$ ) at different joint-level “resolutions”. Specifically, we group the 10 individual joints predicted by COPILLOT into 4 and 6 groups based on semantics (e.g. elbow

and hand joints grouped together), and report the F1-score for each group. This provides multiple evaluation granularities since requirements may differ across various applications. Finally, collision region localization  $\hat{y}^{\text{map}}$  is evaluated in terms of the KL-divergence (*KL*) between predicted and ground truth heatmaps.

##### 5.1.1 Comparison with Baselines

We evaluate our model against existing backbone models and attention schemes for both the collision prediction and localization tasks and summarize the results in Tab. 1. Since we study a novel problem, there is no directly comparable prior work. Hence, we adapt *Time2Col* [38], which performs 2D convolutions on each frame and concatenates features for all frames to perform both collision prediction and localization. Our transformer-based backbone outperforms the convolutional approach of *Time2Col* given a single viewpoint as input (from the root joint, i.e. pelvis).

Under the multi-view setting, we compare two ways of fusing the multi-view information: (1) 3D spatio-temporal attention (*3D ST*) [6], which performs self-attention within each viewpoint, and concatenates features from all viewpoints for prediction, and (2) 4D divided attention over space-time-viewpoint (*4D STV Div*), which extends divided space-time attention [6] by attending only to patches at the same spatial location across viewpoint and time, followed by a spatial self-attention over its own frame. Our proposed space-time-viewpoint attention scheme (*4D STV Full*) outperforms these baselines, due to its larger receptive field and fusing of cross-view information at an earlier stage.

##### 5.1.2 Ablation Study

We ablate important design decisions of COPILLOT and summarize the results in Tab. 2.

**Multi-task learning.** First, we evaluate whether performing both collision prediction and localization in a multi-task fashion improves performance. Models supervised with heatmap prediction are indicated in the *Map* column. As shown, training in this multi-task setting that includes both collision prediction and collision region localization improves prediction performance for both the overall and the per-joint collisions. This trend holds for both single-view and multi-view settings, and the effect is most significant in unseen scenes (right section of Tab. 2). We hypothesize that collision region localization encourages better scene understanding, which helps with cross-scene generalization.

**Multi-view inputs.** We next investigate if using multi-view inputs indeed benefits the prediction task. In our problem formulation, we hypothesize that multiple viewpoints can benefit collision prediction since leveraging all viewpoints can better capture whole-body movement and environment context. This is confirmed by comparing the top section of Tab. 2 (single-view) to the middle section (multi-view). We

Method	Views	Attention	Unseen Motions					Unseen Motions & Scenes				
			Col	F1@4	F1@6	F1@10	KL	Col	F1@4	F1@6	F1@10	KL
Time2Col [38]	Root	-	61.9	56.3	43.9	<b>38.0</b>	0.86	64.3	47.3	33.9	29.4	<b>0.94</b>
COPILOT	Root	3D ST [6]	<b>74.5</b>	<b>57.1</b>	<b>44.7</b>	37.9	<b>0.79</b>	<b>77.1</b>	<b>56.0</b>	<b>40.7</b>	<b>33.2</b>	0.99
COPILOT	All	3D ST [6]	80.9	69.1	55.7	49.6	0.95	<b>84.8</b>	57.4	42.5	37.0	1.41
		4D STV Div	80.0	66.8	57.7	52.3	0.82	81.0	53.7	43.7	37.1	1.25
		4D STV Full (ours)	<b>86.0</b>	<b>71.2</b>	<b>62.2</b>	<b>55.6</b>	<b>0.64</b>	<b>84.8</b>	<b>63.7</b>	<b>52.9</b>	<b>45.6</b>	<b>1.09</b>

Table 1: **Comparison with baselines.** (Top) Our transformer-based backbone outperforms a convolutional baseline [38]. (Bottom) Our proposed 4D attention scheme outperforms alternative approaches to multi-view fusion. All variants perform both collision prediction and localization in a multi-task manner. Classification accuracies and F1 scores are percentages (%).

Input		Output	Unseen M.		Unseen M. & S.	
Views	Modality	Map	Col	F1@10	Col	F1@10
Root	RGB	✗	69.2	<b>38.7</b>	69.9	32.4
Root	RGB	✓	<b>74.5</b>	37.9	<b>77.1</b>	<b>33.2</b>
All	RGB	✗	70.7	43.0	74.7	33.0
All	RGB	✓	<b>86.0</b>	<b>55.6</b>	<b>84.8</b>	<b>45.6</b>
All	Depth	✓	85.0	64.7	87.9	60.1

Table 2: **Ablation study.** Our COPILOT model achieves improved performance when predicting collision heatmaps in addition to classification (multi-task), using multi-view inputs, and using alternative input modalities such as depth.

see a significant performance gain from using multi-view inputs, especially when combined with heatmap prediction, resulting in a 17.7% and 12.4% increase for  $F1@10$  in training and unseen scenes, respectively.

**Multi-modal support.** As COPILOT can be easily extended to incorporate additional sensor modalities, *e.g.* depth maps, we next test our model using depth as input instead of RGB. As expected, this improves performance on both overall and per-joint predictions (bottom of Tab. 2). This is expected because depth provides direct information on scene proximity, which is essential for our task.

## 5.2. Collision Region Localization

Next, we evaluate whether it is necessary to explicitly supervise collision region localization, or if meaningful spatial attention can emerge naturally through the prediction task itself. For this, we consider the single-viewpoint setting, and compare our *COPILOT Root-Only RGB* model to the ablated variant where no collision heatmap is predicted or supervised. We visualize the learned attention weights for this variant using Attention Rollout [1]. Multiple predicted collision region heatmaps are visualized in Fig. 5 (a), where we show three sampled frames for each sequence. We observe that our model with direct heatmap supervision produces meaningful and concentrated heatmaps around reasonable risky regions, unlike the inherent at-

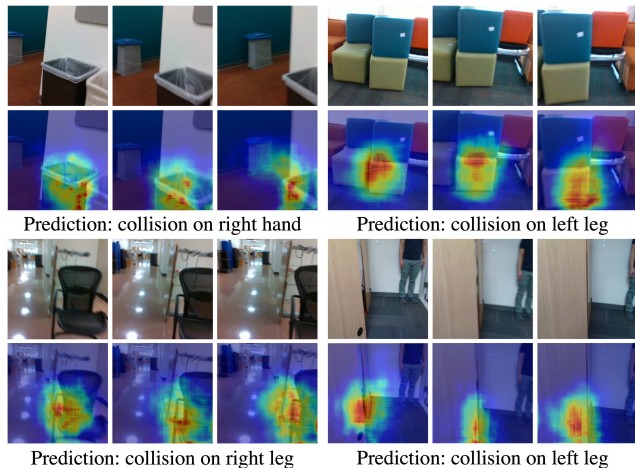


Figure 4: **Sim-to-real transfer.** Predictions from *COPILOT Root-Only RGB* for four real-world videos.

tention learned without direct heatmap supervision. The heatmap offers important information about the potential collision event, *e.g.* in *Scene 2*, the left side of the oven is closer to the person, but only the right side is highlighted since the person moves in that direction.

In Fig. 5 (b), we investigate if the predicted heatmaps are view-consistent. The same timestep is visualized from multiple viewpoints along with the collision region heatmaps from the full *COPILOT RGB* model. The localization is consistent across viewpoints, suggesting that the model has learned a useful representation of the surrounding scene to enable spatial cross-view correspondences.

## 5.3. Generalization to Real-World Videos

We show that our model trained on synthetic data can be directly used on real-world videos (see Fig. 1 & 4). In Fig. 4, we show the results of applying the *COPILOT-Root-Only RGB* model to four real-world videos captured from a chest-mounted camera as a person moves in cluttered offices. Our model reasonably localizes the risky regions that might cause collisions. For the first video (top left), as the person passes the trash bin on the right, our model predicts a potential collision with the right hand and the heatmap highlights

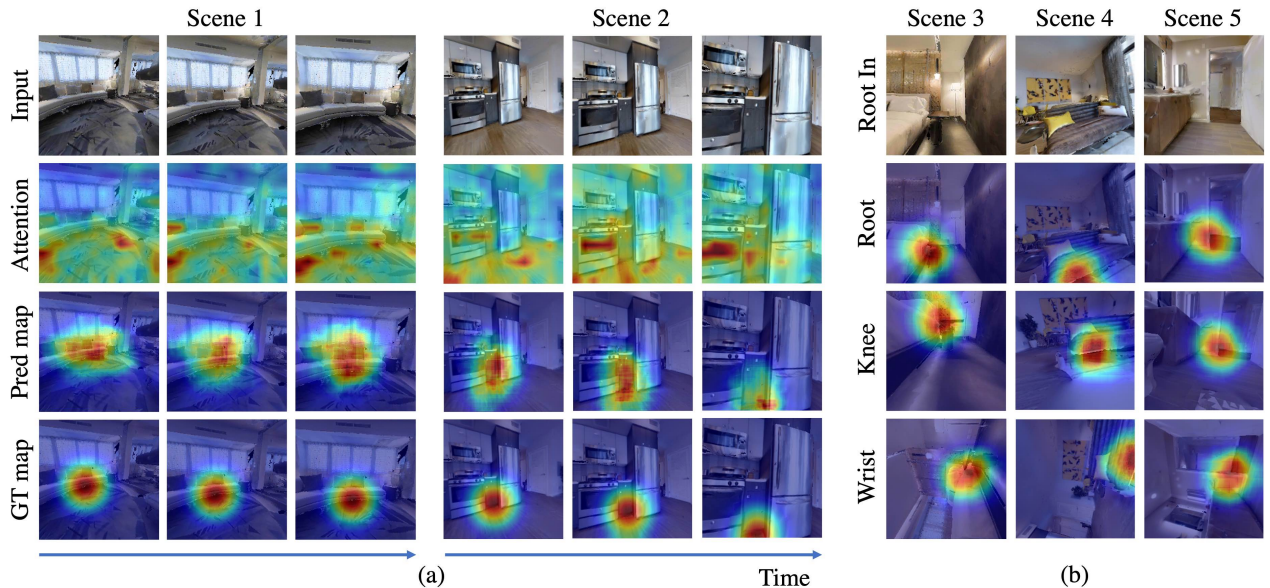


Figure 5: **Collision region localization results on unseen synthetic scenes.** (a) Two examples comparing *COPILOT Root-Only RGB (Pred map)* to the ablated variant that does not predict heatmaps (*Attention*). Explicit heatmap prediction localizes collision much more accurately. (b) Three examples from the full multi-view *COPILOT* model. Predictions from multiple viewpoints are visualized at the same timestep. All examples are on the unseen scenes.

the protruding object. We provide extensive visualizations in the supplementary video demonstrating *COPILOT*'s sim-to-real transfer ability, despite motion blur in the real-world videos and being trained on a different viewpoint.

#### 5.4. Collision Avoidance Assistance

We additionally show that our model's predictions can be used to provide collision avoidance assistance to a virtual human in a simulated environment. Our setup is akin to a hip exoskeleton emulator with actuation to adjust foot placement and center-of-mass states [13]. We approximate the center of mass as the root joint and directly control its orientation. Given the past 1 second of video observations, *COPILOT-Root-RGB* predicts if there will be a collision in the next 1 second. Based on the aggregated probabilities in the left and right halves of the predicted heatmap, the controller slightly rotates the hip by 5 degrees to the left or the right about the z-axis to avoid the higher-probability region. Observations are then sampled at the new rotated pose and the control loop repeats autoregressively.

The policy avoids 35% of collision cases on training scenes and 29% on unseen scenes. Two such cases are visualized in Fig. 6. In the left example, the person originally walks into the vanity, but after the intervention, they instead gradually turn left to avoid the collision.

## 6. Conclusion and Limitations

In this work, we formulate the problem of collision prediction and localization from egocentric videos. To solve




Figure 6: **Collision avoidance visualization.** Grey meshes indicate the history trajectory. Orange is the original future trajectory without intervention, while blue uses collision avoidance assistance. The arrows show the direction of movement.

this problem, we propose a model named *COPILOT* that leverages novel 4D space-time-viewpoint attention to aggregate motion and environment information from multiple viewpoints. We demonstrate that *COPILOT* generalizes to unseen synthetic scenes as well as the real world, and provides useful information for collision avoidance.

Despite the promising results, our model has limitations. Since generating scene-aware 3D human motion sequences is an open research problem [21, 55, 54, 24], the motions in our dataset are scene-agnostic, which provide diverse collisions for training but do not always accurately reflect real-world scenarios. Moreover, our model struggles with multi-modal futures due to ambiguous motion patterns. Finally, developing more sophisticated collision avoidance algorithms that leverage the detailed per-joint outputs of our model is another important direction.



## 7. Acknowledgment

This work was supported by a grant from the Stanford Human-Centered AI Institute (HAI), a Vannevar Bush Faculty Fellowship, an NVIDIA Graduate Fellowship, and the Swiss National Science Foundation under grant number P500PT 206946. We would like to thank colleagues from the Stanford Geometric Computation Group for valuable discussion and constructive feedback.

## References

- [1] Samira Abnar and Willem Zuidema. Quantifying attention flow in transformers. *arXiv preprint arXiv:2005.00928*, 2020. 7
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6836–6846, 2021. 3
- [3] Simi Asher and Kazuo Kiguchi. Real-time accident prediction using deep learning for. In *2020 IEEE International Conference on Real-time Computing and Robotics (RCAR)*, pages 333–338, 2020. 2, 3
- [4] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 5
- [5] Wenkai Bao, Dario Villarreal, and J-C Chiao. Vision-based autonomous walking in a lower-limb powered exoskeleton. In *2020 IEEE 20th International Conference on Bioinformatics and Bioengineering (BIBE)*, pages 830–834. IEEE, 2020. 2, 3
- [6] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding. *arXiv preprint arXiv:2102.05095*, 2(3):4, 2021. 3, 4, 5, 6, 7
- [7] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 3
- [8] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*, 2017. 2, 3
- [9] Yuchen Chang, Keigo Matsumoto, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. Redirection controller using reinforcement learning. *IEEE Access*, 9:145083–145097, 2021. 2
- [10] Bing Chen, Chun-Hao Zhong, Xuan Zhao, Hao Ma, Xiao Guan, Xi Li, Feng-Yan Liang, Jack Chun Yiu Cheng, Ling Qin, Sheung-Wai Law, et al. A wearable exoskeleton suit for motion assistance to paralysed patients. *Journal of orthopaedic translation*, 11:7–18, 2017. 1
- [11] Chuheng Chen, Kuangen Zhang, Yuquan Leng, Xinxing Chen, and Chenglong Fu. Unsupervised sim-to-real adaptation for environmental recognition in assistive walking. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2022. 2
- [12] Lung-Pan Cheng, Eyal Ofek, Christian Holz, and Andrew D Wilson. Vroamer: generating on-the-fly vr experiences while walking inside large, unknown real-world building environments. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 359–366. IEEE, 2019. 1, 2
- [13] Vincent L Chiu, Michael Raitor, and Steven H Collins. Design of a hip exoskeleton with actuation in frontal and sagittal planes. *IEEE Transactions on Medical Robotics and Bionics*, 3(3):773–782, 2021. 8
- [14] Erwin Coumans and Yunfei Bai. Pybullet, a python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>, 2016–2021. 3
- [15] Tianyang Dong, Xianwei Chen, Yifan Song, Wenyuan Ying, and Jing Fan. Dynamic artificial potential fields for multi-user redirected walking. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 146–154. IEEE, 2020. 1, 2
- [16] 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. *arXiv preprint arXiv:2010.11929*, 2020. 3
- [17] William Falcon et al. Pytorch lightning. *GitHub. Note: https://github.com/PyTorchLightning/pytorch-lightning*, 3, 2019. 6
- [18] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6824–6835, 2021. 3
- [19] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019. 3
- [20] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1933–1941, 2016. 3
- [21] Mohamed Hassan, Duygu Ceylan, Ruben Villegas, Jun Saito, Jimei Yang, Yi Zhou, and Michael J Black. Stochastic scene-aware motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11374–11384, 2021. 8
- [22] Christian Hirt, Markus Zank, and Andreas Kunz. Geometry extraction for ad hoc redirected walking using a slam device. In *Augmented Reality, Virtual Reality, and Computer Graphics: 5th International Conference, AVR 2018, Otranto, Italy, June 24–27, 2018, Proceedings, Part I 5*, pages 35–53. Springer, 2018. 1, 2
- [23] Jeannette E Holm. Collision prediction and prevention in a simultaneous multi-user immersive virtual environment. Master’s thesis, Miami University, 2012. 1
- [24] Siyuan Huang, Zan Wang, Puhao Li, Baoxiong Jia, Tengyu Liu, Yixin Zhu, Wei Liang, and Song-Chun Zhu. Diffusion-based generation, optimization, and planning in 3d scenes. *arXiv preprint arXiv:2301.06015*, 2023. 8

- [25] Kübra Karacan, Jan T Meyer, H İşil Bozma, Roger Gassert, and Evren Samur. An environment recognition and parameterization system for shared-control of a powered lower-limb exoskeleton. In *2020 8th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechanics (BioRob)*, pages 623–628. IEEE, 2020. 2, 3
- [26] Kazuo Kiguchi and Ryosuke Matsuo. Accident prediction based on motion data for perception-assist with a power-assist robot. In *2017 IEEE symposium series on computational intelligence (SSCI)*, pages 1–5. IEEE, 2017. 2
- [27] Andrew Garrett Kurbis, Brokoslaw Laschowski, and Alex Mihailidis. Stair recognition for robotic exoskeleton control using computer vision and deep learning. *bioRxiv*, 2022. 2, 3
- [28] Brokoslaw Laschowski, William McNally, Alexander Wong, and John McPhee. Environment classification for robotic leg prostheses and exoskeletons using deep convolutional neural networks. *Frontiers in Neurorobotics*, 15:730965, 2022. 2, 3
- [29] Dong-Yong Lee, Yong-Hun Cho, and In-Kwon Lee. Real-time optimal planning for redirected walking using deep q-learning. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 63–71. IEEE, 2019. 2
- [30] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 5
- [31] Du-Xin Liu, Jing Xu, Chunjie Chen, Xingguo Long, Dacheng Tao, and Xinyu Wu. Vision-assisted autonomous lower-limb exoskeleton robot. *IEEE transactions on systems, man, and cybernetics: systems*, 51(6):3759–3770, 2019. 2
- [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 *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 704–721. Springer, 2020. 3
- [33] Shaowei Liu, Subarna Tripathi, Somdeb Majumdar, and Xiaolong Wang. Joint hand motion and interaction hotspots prediction from egocentric videos. *arXiv preprint arXiv:2204.01696*, 2022. 3
- [34] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. *arXiv preprint arXiv:2106.13230*, 2021. 3
- [35] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015. 3
- [36] Zhengyi Luo, Ryo Hachiuma, Ye Yuan, and Kris Kitani. Dynamics-regulated kinematic policy for egocentric pose estimation. *Advances in Neural Information Processing Systems*, 34:25019–25032, 2021. 3
- [37] Karttikeya Mangalam, Haoqi Fan, Yanghao Li, Chao-Yuan Wu, Bo Xiong, Christoph Feichtenhofer, and Jitendra Malik. Reversible vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10830–10840, 2022. 3
- [38] Aashi Manglik, Xinshuo Weng, Eshed Ohn-Bar, and Kris M Kitani. Forecasting time-to-collision from monocular video: Feasibility, dataset, and challenges. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8081–8088. IEEE, 2019. 3, 6, 7
- [39] Evonne Ng, Donglai Xiang, Hanbyul Joo, and Kristen Grauman. You2me: Inferring body pose in egocentric video via first and second person interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9890–9900, 2020. 3
- [40] Jiangpeng Ni, Chunjie Chen, Zhuo Wang, Youfu Liu, Xinyu Wu, and Yida Liu. A control method with terrain classification and recognition for lower limb soft exosuit. In *2020 IEEE International Conference on E-health Networking, Application & Services (HEALTHCOM)*, pages 1–6. IEEE, 2021. 2
- [41] Yuepeng Qian, Yining Wang, Chuheng Chen, Jingfeng Xiong, Yuquan Leng, Haoyong Yu, and Chenglong Fu. Predictive locomotion mode recognition and accurate gait phase estimation for hip exoskeleton on various terrains. *IEEE Robotics and Automation Letters*, 2022. 2
- [42] Sharif Razzaque, Zachariah Kohn, and Mary C. Whitton. Redirected Walking. In *Eurographics 2001 - Short Presentations*. Eurographics Association, 2001. 1, 2
- [43] Davis Rempe, Tolga Birdal, Aaron Hertzmann, Jimei Yang, Srinath Sridhar, and Leonidas J Guibas. Humor: 3d human motion model for robust pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11488–11499, 2021. 2, 3
- [44] Sebastian Ruder. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747*, 2016. 6
- [45] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9339–9347, 2019. 2, 4
- [46] Abhishek Sharma and Eric Rombokas. Improving imu-based prediction of lower limb kinematics in natural environments using egocentric optical flow. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 30:699–708, 2022. 2, 3
- [47] Andrew Szot, Alexander Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Singh Chiplot, Oleksandr Maksymets, et al. Habitat 2.0: Training home assistants to rearrange their habitat. *Advances in Neural Information Processing Systems*, 34, 2021. 2, 4
- [48] Jerald Thomas and Evan Suma Rosenberg. A general reactive algorithm for redirected walking using artificial potential functions. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 56–62. IEEE, 2019. 1, 2
- [49] Denis Tome, Patrick Peluse, Lourdes Agapito, and Hernan Badino. xr-egopose: Egocentric 3d human pose from an hmd camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7728–7738, 2019. 3
- [50] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with

- 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. 3
- [51] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3
- [52] Conor James Walsh, Ken Endo, and Hugh Herr. A quasi-passive leg exoskeleton for load-carrying augmentation. *International Journal of Humanoid Robotics*, 4(03):487–506, 2007. 1, 3
- [53] Jian Wang, Lingjie Liu, Weipeng Xu, Kripasindhu Sarkar, and Christian Theobalt. Estimating egocentric 3d human pose in global space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11500–11509, 2021. 3
- [54] Jingbo Wang, Yu Rong, Jingyuan Liu, Sijie Yan, Dahua Lin, and Bo Dai. Towards diverse and natural scene-aware 3d human motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20460–20469, 2022. 8
- [55] Jiashun Wang, Huazhe Xu, Jingwei Xu, Sifei Liu, and Xiaolong Wang. Synthesizing long-term 3d human motion and interaction in 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9401–9411, 2021. 8
- [56] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. 3
- [57] Niall L Williams, Aniket Bera, and Dinesh Manocha. Redirected walking in static and dynamic scenes using visibility polygons. *IEEE transactions on visualization and computer graphics*, 27(11):4267–4277, 2021. 1
- [58] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 284–293, 2019. 3
- [59] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. 5
- [60] Fei Xia, Amir R. Zamir, Zhi-Yang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: real-world perception for embodied agents. In *Computer Vision and Pattern Recognition (CVPR), 2018 IEEE Conference on*. IEEE, 2018. 2, 3
- [61] Ye Yuan and Kris Kitani. Ego-pose estimation and forecasting as real-time pd control. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10082–10092, 2019. 3