

A. Overview

In this supplementary material, we provide additional details on the dataset (Appendix B), experiments (Appendix C), generated videos (Appendix D), and an overall algorithm (Appendix E) of ZeroHSI. We highly recommend viewing our [project page](#) for compelling demonstrations across diverse scenarios.

B. AnyInteraction Dataset

In this section, we elaborate on the statistics of the AnyInteraction dataset. We summarize these statistics in Tab. S1 and visualize our AnyInteraction dataset in Fig. S1.

B.1. Scenes

Our AnyInteraction dataset consists of diverse scenes from TRUMANS dataset [34], public 3D assets libraries, and reconstructed real scenes from the Mip-NeRF 360 dataset [4] and Tanks and Temples dataset [41], resulting in 7 indoor scenes (*Bedroom, Living Room, Gym, Bar, Greenhouse, Store, Room*) and 5 outdoor scenes (*Playground, Cafe, Garden, Bicycle, Truck*).

Among our synthetic scenes, *Playground* and *Cafe* are manually composed using models from 3D asset libraries, while the remaining six scenes are sourced directly from asset libraries with adjustments to their scale and layout to facilitate interactions. For 3DGS [38] reconstruction, we generate 300-500 cameras per scene and manually filter occluded and low-quality views. We then render RGBD images to obtain initial point clouds and perform the reconstruction using the official 3D Gaussian Splatting implementation [38].

For real scenes, we reconstruct them using images and camera views from the official datasets. We scale and transform these scenes to align their ground level and match real-world sizes. For visualization, we further extract scene meshes using SuGaR [19].

B.2. Dynamic Objects

AnyInteraction includes 7 types of dynamic objects (*Guitar, Barbell, Watering Can, Office Chair, Shopping Cart, Vase, Mower*), and all of them are rigid. We obtain these objects from public 3D Assets libraries. Similar to the reconstruction process of the synthetic scenes, we generate 70 cameras per object, render RGBD images to obtain initial point clouds, and reconstruct using the official 3D Gaussian Splatting implementation [38].

B.3. Evaluation Instances

As shown in Tab. S1 and Fig. S1, our AnyInteraction dataset contains 22 evaluation instances including 13 static interaction instances and 9 dynamic object interaction instances. Each interaction instance comprises a text prompt and initial state. The initial state typically features a standing pose with

nearby objects. We adjust the standing pose for specific instances to ease the interaction video generation.

C. Additional Experiment Details and Results

In this section, we provide additional; introduction to the experimental settings (Appendix C.1), implementation (Appendices C.2 to C.4), human studies (Appendix C.5), and baseline comparisons (Appendix C.6). We explain how our method is capable of synthesizing HSIs of varying lengths and present additional qualitative results of diverse long-term human-scene interactions (Appendix C.7). We show qualitative examples of our ablation study (Appendix C.8). We also demonstrate our approach’s flexibility by integrating various state-of-the-art video generation models beyond the model used in our main experiments (Appendix C.9).

C.1. Experimental Settings

For the static scenarios, we evaluate on 11 static instances (excluding *Bicycle* and *Truck* scenes). We evaluate each scene with 5 different seeds (as we consider a generative setting), yielding 55 generated motion sequences for evaluation. For the dynamic scenarios, we evaluate on all 9 dynamic object interaction instances. Similarly, we evaluate each scene with 5 different seeds and this leads to 45 motion sequences.

C.2. Implementation Details

We employ KLING image-to-video model v1.0 [76] for HSI video generation and SAM 2.0 [69] for segmentation. For efficiency, we downsample the generated 153-frame videos to 51 frames for all experiments. For initialization, we employ a default standing pose with objects positioned near the human. We make adjustments for the initial pose and camera view for some instances. For comparative evaluation, we utilize the official implementations and pre-trained models of the baseline methods. Since all baseline approaches require spatial conditions for navigation, we provide them with the human/object trajectories generated by our method. We also supply the initialization configuration to these baselines, as they either require or can accommodate initial-state inputs.

The Adam optimizer [40] is utilized across all optimization stages: camera pose estimation, 4D HSI optimization, color net fine-tuning, and refinement. For camera pose estimation, we optimize 30 iterations per frame with a learning rate of 0.001. In 4D HSI optimization, we perform 300 iterations per frame. During the initial 30 iterations, the optimization of human poses is limited to root translation and global orientation. We use a learning rate of 0.01, and the loss weights are set as $\lambda = 0.1$, $\lambda_{\text{center}} = 0.001$, and $\lambda_{\text{depth}} = 0.001$. The color net is fine-tuned simultaneously every 5 step during 4D HSI optimization with a learning rate of 0.00001. For refinement, we optimize 1000 iterations for the entire sequence with a learning rate of 0.05 and physics loss weight $\lambda_{\text{physics}} = 0.001$.

Scene Name	Scene Type	Source	Text Prompt	Objects
Bedroom	Indoor	TRUMANS [34]	<i>The person is sitting on the bed.</i> <i>The person is sitting on the windowsill.</i> <i>The person is leaning on the ladder.</i>	Static Static Static
Living Room	Indoor	TRUMANS [34]	<i>The person is sitting on the table.</i> <i>The person is sitting on the sofa back.</i> <i>The person is playing guitar while sitting on the sofa.</i>	Static Static Guitar
Gym	Indoor	Asset Libraries	<i>The person is running on the treadmill.</i> <i>The person is lifting weights.</i>	Static Barbell
Bar	Indoor	Asset Libraries	<i>The person is leaning on the bar.</i>	Static
Playground	Outdoor	Asset Libraries	<i>The person is sliding down the slide.</i>	Static
Greenhouse	Indoor	Asset Libraries	<i>The person is sitting on the chair.</i> <i>The person is watering flowers with a watering can.</i>	Static Watering Can
Cafe	Outdoor	Asset Libraries	<i>The person is sitting on the chair.</i>	Static
Store	Indoor	Asset Libraries	<i>The person is picking out snacks on the shelf.</i> <i>The person is sliding while sitting on the chair.</i> <i>The person is pushing shopping cart.</i>	Static Office Chair Shopping Cart
Garden	Outdoor	Mip-NeRF 360 [4]	<i>The person is watering flowers with a watering can.</i> <i>The person is lifting a vase.</i> <i>The person is operating a lawn mower.</i>	Watering Can Vase Mower
Bicycle	Outdoor	Mip-NeRF 360 [4]	<i>The person is sitting on the bench.</i>	Static
Room	Indoor	Mip-NeRF 360 [4]	<i>The person is playing guitar while sitting on the sofa.</i>	Guitar
Truck	Outdoor	Tanks&Temples [41]	<i>The person is cleaning the car.</i>	Static

Table S1. Statistics of our AnyInteraction dataset.

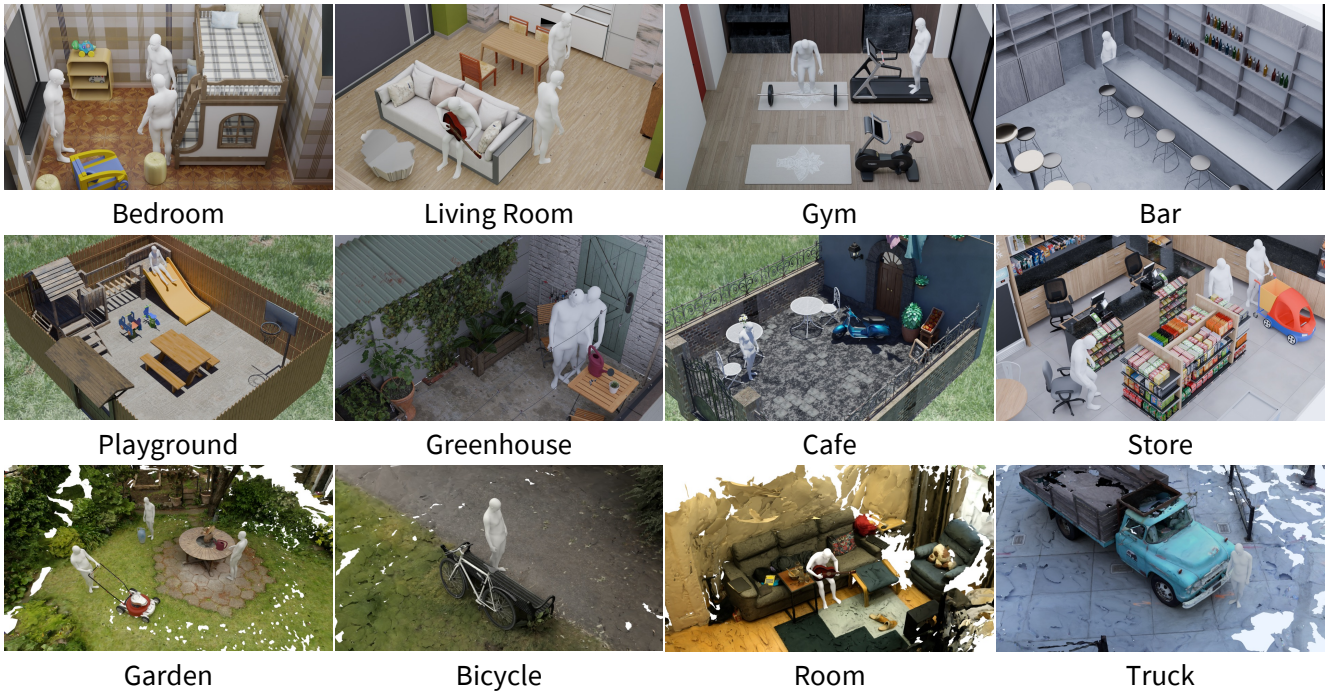


Figure S1. Visualization of our AnyInteraction dataset.

C.3. Details on Initial Camera View Selection

Based on the assumption that camera movement remains minimal over brief durations (approximately 5 seconds), we address occlusion challenges by selecting an initial camera view \mathbf{T}_0 that maximizes visibility of both human subjects and dynamic objects in the first frame. We select camera from a predefined camera array distributed across concentric spherical surfaces with varying radii, all oriented toward the same sphere center.

The sphere center’s height is set as $c^Z = h_{\text{pelvis}} + 0.1$ in a z-up coordinate system. For static scenes, its horizontal position equals the pelvis projection $\mathbf{p}_{\text{pelvis}}^{\text{XoY}}$. For dynamic scenes, we identify the human’s orientation direction \mathbf{d} and create a ray l from $\mathbf{p}_{\text{pelvis}}^{\text{XoY}}$ along \mathbf{d} . We then project dynamic object particles onto l and set c^{XoY} as the midpoint between $\mathbf{p}_{\text{pelvis}}^{\text{XoY}}$ and the furthest projection.

For camera array generation, we place cameras on spherical surfaces at different radii, using only the hemisphere facing the human but avoiding direct frontal views. All cameras are oriented toward the sphere center. This setup balances complete human capture while reducing depth ambiguity in the generated video.

For each candidate camera, we render two depth maps: D (entire scene) and $D_{\mathcal{H}}$ (human only). Non-occluded areas are defined as $\Omega = \{(u, v) \mid |D(u, v) - D_{\mathcal{H}}(u, v)| < \epsilon\}$. We project the first 22 SMPL-X joints [60] onto the image plane and count visible joints as N_{visible} . The camera with maximum N_{visible} is selected. With equal visibility, we choose the camera closest to the equatorial plane for better accuracy for height reconstruction.

C.4. Details on Physics Loss

We leverage a hand-object contact loss to encourage hand-object contact when they are in close proximity after the optimization process described in Sec. 3.4. We first define the contact set \mathcal{C}_t for each frame as:

$$\mathcal{C}_t = \left\{ (\mathbf{v}_i, \boldsymbol{\mu}_p) \mid \|\mathbf{v}_i - \boldsymbol{\mu}_p\|_2 < \epsilon, \mathbf{v}_i \in \mathbf{H}_t, \boldsymbol{\mu}_p \in \boldsymbol{\mu}_{\mathcal{O}}^t \right\}, \quad (\text{S1})$$

where \mathbf{H}_t represents SMPL-X [60] hand vertices and $\boldsymbol{\mu}_{\mathcal{O}}^t$ denotes Gaussian object particle positions in frame t . The frame-wise hand-object contact loss is then defined as:

$$\mathcal{L}_{\text{contact}}^t = \frac{1}{\#\mathcal{C}_t} \min_{\mathbf{d}} \left\{ \mathbf{d} = \|\mathbf{v}_i - \boldsymbol{\mu}_p\|_2 \mid (\mathbf{v}_i, \boldsymbol{\mu}_p) \in \mathcal{C}_t \right\}, \quad (\text{S2})$$

which we only apply when $\#\mathcal{C}_t > 0$.

Additionally, since VPoser [60] as a human pose prior does not inherently ensure sequence smoothness after refinement, we apply an additional smoothness loss between adjacent frames:

$$\mathcal{L}_{\text{smooth}} = \frac{1}{T-1} \sum_{t=0}^{T-1} \|\boldsymbol{\Theta}_t - \boldsymbol{\Theta}_{t+1}\|_2. \quad (\text{S3})$$

Survey Questions

In the following section, you will watch 10 video pairs (each pair consists of a left video and a right video).

Each video pair will be shown once, and then you will be asked two questions about it.

Pair 1: Compare the two videos below. Which one has more realistic human motion? *



- ☐ Left video shows more realistic human motion
- ☐ Right video shows more realistic human motion

Pair 1: See the videos above again. Which one looks more like "playing guitar while sitting on the sofa"? *

- ☐ Left video looks more like playing guitar while sitting on the sofa
- ☐ Right video looks more like playing guitar while sitting on the sofa

Pair 2: Compare the two videos below. Which one has more realistic human motion? *



Figure S2. Screenshot of the human study interface.

The overall physics loss is defined as:

$$\mathcal{L}_{\text{physics}} = \lambda_{\text{contact}} \frac{1}{T} \sum_{t=0}^T \mathcal{L}_{\text{contact}}^t + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}}, \quad (\text{S4})$$

where we set $\lambda_{\text{contact}} = 1$ and use $\lambda_{\text{smooth}} = 0.3$ for static scenarios and $\lambda_{\text{smooth}} = 0.1$ for dynamic scenarios.

C.5. Human Study Details

We recruit 400 participants via the Prolific platform. The participants are divided into 4 groups, each of which includes 100 participants. Each participant is shown two side-by-side videos and forced to choose one from them. One of the videos is ours, and the other is generated by a baseline method. The left-right order is randomized. Each participant



(a) “Walking forward to the table.”



(b) “Watering flowers with a watering can.”



(c) “Putting down the watering can.”

Figure S3. **Example of long-term human-scene interaction sequence.** ZeroHSI generates this long interaction sequence from three 5-second HSI videos.

is shown 10 pairs of videos. For each pair of videos, the participant is asked two questions: (1) to choose the video that is higher quality, and (2) to choose the video that is better aligned with the text prompt. We show a screenshot of the human study interface in Fig. S2.

C.6. Comparison Details

In each generation process, our ZeroHSI outputs a 5-second HSI sequence at 10 fps. All baseline methods output sequences at 30 fps, which we downsample to 10 fps for comparison. For a fair comparison, we do not apply refinement




Figure S4. **Example of long-term human-scene interaction sequence.** ZeroHSI generates this long interaction sequence from two 5-second HSI videos (sitting down, standing up) and one 10-second HSI video (practicing Tai Chi).

to our outputs or baseline results. Since TRUMANS [34] and LINGO [33] require occupancy grid inputs, we convert meshes from both synthetic scenes and real scenes (via SuGaR [19]) into occupancy grids, yielding water-tight meshes as a byproduct. We use these water-tight meshes for penetration metric calculations, as they enable proper inside-outside definition necessary for Signed Distance Field (SDF) computation.

C.7. Long-Term Interaction Synthesis

As shown in Fig. 1 (first walking forward, then watering flowers), our method inherently supports the generation of long-term HSI sequences. While KLING image-to-video mode is limited to generating videos under 10 seconds and only allows 5-second extensions per generation, we overcome this limitation by using the last frame of each generated HSI sequence as the initial state for the subsequent generation process. By repeating the entire generation process described in Sec. 3, we can synthesize longer sequences. This approach ensures consistent quality across each gener-

ated sequence while enabling flexible control over individual video clips. For dynamic object interactions, the depth regularization term varies between different 5-second video clips, making the constant depth assumption within each sequence reasonable.

Fig. S3 presents a 15-second interaction sequence where the character walks forward, waters flowers, and puts down the watering can. We achieve this extended interaction by generating and reconstructing three HSI videos into a cohesive 4D sequence.

Fig. S4 and Fig. S5 showcase another 20-second sequence where the character sits on a bench, stands up, and performs Tai Chi. Leveraging diverse video generation capabilities, ZeroHSI can create multiple variations of this interaction, as demonstrated by the alternative Tai Chi sequences (ii) and (iii) in Fig. S5.

C.8. Qualitative Results of Ablations

We show qualitative results of our ablation study on our per-frame optimization approach. The examples in Fig. S6




Figure S5. **Example of long-term human-scene interaction sequence.** ZeroHSI generates this long interaction sequence from two 5-second HSI videos (sitting down, standing up) and one 10-second HSI video (practicing Tai Chi). We show multiple variations of practicing Tai Chi generated by ZeroHSI, demonstrating its diverse generation capability.

reveal that both WHAM [72] and WHAC [93] fail to estimate correct global translation, resulting in penetrations between human and scene. Meanwhile, rendered results across three consecutive frames show our optimization-based method achieving smoother motions compared to "ZeroHSI w/o

OPT_{body} ". These results further demonstrate the effectiveness and technical contribution of our per-frame optimization approach on reconstructing 4D HSIs from 2D videos.

We present additional results in Fig. S7 by rendering the reconstructed interactions through applying SMPL-X pa-




Figure S6. **Qualitative results of ablation study on our optimization-based HSI motion reconstruction.** Our full method reconstructs root translation more accurately than WHAM [72] and WHAC [93] while achieving smoother results than ZeroHSI w/o OPT_{body}.



Figure S7. **Rendered ablation study results.** We render the reconstructed interactions by applying SMPL-X parameters to Gaussian avatars. WHAM [72] and WHAC [93] show abrupt changes between rendered frames due to human-scene penetrations.



Figure S8. **HSI generated with KLING 1.5 [76].** Text prompt: “Sitting on the table.”

rameters to Gaussian avatars and visualizing via Gaussian rasterization. Due to the lack of scene awareness in both WHAM and WHAC, although the reconstructed motion sequences appear smooth, the rendered results show abrupt

changes between adjacent frames as the human progressively penetrates the scene geometry. These artifacts correspond to the relatively low CLIP scores and CLIP consistency reported in Tab. 4.



Figure S9. HSI generated with KLING 1.6 [76]. Text prompt: “Watering flowers with a watering can.”



Figure S10. HSI generated with Wan 2.1 [77]. Text prompt: “Picking out snacks on the shelf.”





Figure S11. HSI generated with Sora [9].

C.9. Results with Other Video Generation Models

We use off-the-shelf video generation model to ensure ZeroHSI maintains compatibility with various existing tech-

niques and remains adaptable to future advancements in video generation. To evaluate ZeroHSI’s adaptability across different models, we test it with other state-of-the-art image-to-video models: KLING 1.5 [76], KLING 1.6 [76], Wan



(a) A **cloth** appears in hand during car wiping; a **mailbox** appears behind as the character leans forward. (b) The **bicycle seat** appears white in the generated video while it is black in the real scene.

Figure S12. **Example of incorrect contents.** The first row shows frames containing unwanted contents. The second row shows results using only static background masks, while the third row shows our aggregated dynamic foreground masking approach. The fourth row presents our final reconstruction results.



(a) As the character squats down, their **right hand** fades and blends into the background shelf. (b) As the character walks forward, their **left hand** fades and blends into the background cabinet.

Figure S13. **Example of body part disappearance.** The first row shows frames of body part disappearance. The second row presents our final reconstruction results.

2.1 [77], and Sora [9].

Qualitative results in Figs. S8 to S11 demonstrate Ze-

roHSI's effectiveness across different open-source video generation models and proprietary models with external




Figure S14. **Example of human appearance change.** The first row shows frames of human appearance change. The character’s **skin tone and hairstyle** change while leaning forward; **black patterns** incorrectly shift from the sleeves onto the torso. The second row shows gaussian rendering results after optimization. The third row presents our final reconstruction results.

APIs. ZeroHSI performs well with newer versions of KLING in both static (“sitting on table”) and dynamic (“watering flowers”) scenarios. With Wan 2.1, ZeroHSI successfully generates plausible interaction motions that follow the provided text prompt. We observe that Sora struggle to generate videos aligning well with the given text. However, when we examine it without textual input, our ZeroHSI still generates meaningful 4D HSI of “watching the vase on the table.”

D. Example of Low-Quality Video

In this section, we demonstrate the robustness of our method by presenting challenging examples of low-quality video inputs. Specifically, we incorporate techniques for handling incorrect content in camera pose estimation, as well as dealing with body part disappearance and human appearance changes during HSI reconstruction. For each issue, we provide representative examples and show how our method successfully addresses these challenges through the reconstruction results.

D.1. Incorrect Contents

Video generation models produce incorrect contents in regions initially occluded by humans or objects that later become visible, and we use the aggregating dynamic foreground masks from frames 0 to t to eliminate incorrect contents. Fig. S12 presents two examples of such incorrect generated content. In the first example, a cloth appears in the character’s hand during car wiping, and a mailbox appears

behind the character as they lean forward. In the second example, the color of the bicycle seat behind the human is generated incorrectly (black in the real scene but white in the generated video).

We compare two approaches: directly using the static background mask in the current frame (second row) versus using the aggregating dynamic foreground masks from frames 0 to t (third row). Using only the static background mask in the current frame fails to mask out all unwanted contents, while using the complementary set of the aggregated dynamic foreground mask successfully eliminates all incorrect contents. The plausible reconstruction results in the last row demonstrate the effectiveness of our design.

D.2. Body Part Disappearance

Fig. S13 demonstrates body part disappearance in both static and dynamic scenarios, where characters’ hands gradually fade and blend into the background. Our per-frame optimization method produces natural and cohesive reconstructed HSI (second row), demonstrating its robustness to such issues.

D.3. Change of Human Appearance

Video generation models often produce inconsistent human appearance. To address this, we fine-tune the color net defined in Eq. (5) during optimization using the photometric loss in each frame (Eq. (10)). Fig. S14 shows an example where a character’s skin tone, hairstyle, and clothing patterns

change incorrectly. The second row displays gaussian rendering results after per-frame optimization, demonstrating how our approach successfully corrects appearance inconsistencies by gradually fine-tuning the color net. The final row shows plausible reconstruction results, validating our design’s effectiveness.

E. Algorithms

We summarize ZeroHSI and show the overall algorithm in Alg. 1.

Algorithm 1 ZeroHSI: Zero-shot Human-Scene Interaction Generation

```

1: Input:
2: Scene Gaussians  $\mathcal{G}_S$ , Object Gaussians  $\mathcal{G}_O$ 
3: Initial human pose  $\mathcal{M}_0 = (\mathbf{r}_0, \phi_0, \Theta_0)$ 
4: Initial object pose  $\mathbf{P}_0$ , Initial camera pose  $\mathbf{T}_0$ 
5: Text prompt  $c$  describing the interaction
6: function GENERATEHSIVIDEO
7:    $\mathcal{G}_H^0 \leftarrow \mathcal{A}(\mathbf{r}_0, \phi_0, \Theta_0; \mathbf{T}_0)$  ▷ Initialize human Gaussians
8:    $\mathcal{G}_O^0 \leftarrow \mathcal{G}_O(\mathbf{P}_0)$  ▷ Transform object Gaussians
9:    $\mathbf{I}_0 \leftarrow \mathcal{R}(\mathcal{G}_H^0, \mathcal{G}_O^0, \mathcal{G}_S; \mathbf{T}_0)$  ▷ Render initial frame
10:   $\{\mathbf{I}_t\}_{t=0}^T \leftarrow \text{VideoGen}(\mathbf{I}_0, c)$  ▷ We use KLING
11:   $\{\mathbf{M}_H^t, \mathbf{M}_O^t\}_{t=0}^T \leftarrow \text{SAM2}(\{\mathbf{I}_t\}_{t=0}^T)$  ▷ Segment video
12:  return  $\{\mathbf{I}_t, \mathbf{M}_H^t, \mathbf{M}_O^t\}_{t=0}^T$ 
13: end function
14: function RECONSTRUCT4DHSI
15:  for  $t = 1$  to  $T$  do
16:     $\mathbf{T}_* \leftarrow \arg \min_{\mathbf{T}} \mathcal{L}_2(\mathcal{R}(\mathcal{G}_S(\mathbf{T}); \mathbf{T}_{t-1}) \odot \mathbf{M}_t, \mathbf{I}_t \odot \mathbf{M}_t)$  ▷ Estimate camera pose
17:     $\mathbf{T}_t \leftarrow \mathbf{T}_*^{-1} \mathbf{T}_{t-1}$ 
18:    Initialize  $\mathbf{r}_t, \phi_t$  from previous frame ▷ Initialize global transform
19:    Initialize  $\Theta_t$  from pose estimation model
20:     $\mathcal{M}_t, \mathbf{P}_t \leftarrow \arg \min(\mathcal{L}_{\text{rgb}} + \lambda_{\text{center}} \mathcal{L}_{\text{center}} + \lambda_{\text{depth}} \mathcal{L}_{\text{depth}})$  ▷ Optimize human and object poses
21:  end for
22:  return  $\{(\mathcal{M}_t, \mathbf{P}_t)\}_{t=1}^T$ 
23: end function
24: function REFINEMENT
25:  for  $t = 1$  to  $T$  do
26:     $\hat{J}_t \leftarrow \text{SMPL-X}(\mathcal{M}_t)$  ▷ Get reference joints
27:     $\mathcal{L}_{\text{fit}}^t \leftarrow \mathcal{L}_2(\hat{J}_t, J_t(\mathbf{r}_t, \phi_t, \mathcal{D}(\mathbf{z}_t)))$  ▷ Compute fitting loss for frame  $t$ 
28:  end for
29:   $\mathcal{L}_{\text{physics}} \leftarrow \text{CalculatePhysicsLoss}(\{\mathbf{r}_t, \phi_t, \Theta_t\}_{t=1}^T)$  ▷ Compute physics loss
30:   $\mathbf{r}, \phi, \mathbf{z} \leftarrow \arg \min(\frac{1}{T} \sum_{t=1}^T \mathcal{L}_{\text{fit}}^t + \lambda_{\text{physics}} \mathcal{L}_{\text{physics}})$ 
31:  for  $t = 1$  to  $T$  do
32:     $\Theta_t \leftarrow \mathcal{D}(\mathbf{z}_t)$  ▷ Decode VPoser latent
33:     $\mathcal{M}_t \leftarrow (\mathbf{r}_t, \phi_t, \Theta_t)$  ▷ Update human pose
34:  end for
35:  return  $\{(\mathcal{M}_t, \mathbf{P}_t)\}_{t=1}^T$ 
36: end function
37:  $\{\mathbf{I}_t, \mathbf{M}_H^t, \mathbf{M}_O^t\}_{t=0}^T \leftarrow \text{GenerateHSIVideo}()$ 
38:  $\tau \leftarrow \text{Reconstruct4DHSI}()$ 
39:  $\tau \leftarrow \text{Refinement}()$ 
40: Output: 4D HSI sequence  $\tau = \{(\mathcal{M}_t, \mathbf{P}_t)\}_{t=1}^T$ 

```
