

Pose2Room: Understanding 3D Scenes from Human Activities

Yinyu Nie¹, Angela Dai¹, Xiaoguang Han², and Matthias Nießner¹

¹ Technical University of Munich

² The Chinese University of Hong Kong, Shenzhen

Abstract. With wearable IMU sensors, one can estimate human poses from wearable devices without requiring visual input. In this work, we pose the question: Can we reason about object structure in real-world environments solely from human trajectory information? Crucially, we observe that human motion and interactions tend to give strong information about the objects in a scene – for instance a person sitting indicates the likely presence of a chair or sofa. To this end, we propose P2R-Net to learn a probabilistic 3D model of the objects in a scene characterized by their class categories and oriented 3D bounding boxes, based on an input observed human trajectory in the environment. P2R-Net models the probability distribution of object class as well as a deep Gaussian mixture model for object boxes, enabling sampling of multiple, diverse, likely modes of object configurations from an observed human trajectory. In our experiments we show that P2R-Net can effectively learn multi-modal distributions of likely objects for human motions, and produce a variety of plausible object structures of the environment, even without any visual information. The results demonstrate that P2R-Net consistently outperforms the baselines on the PROX dataset and the VirtualHome platform.

Keywords: 3D Scene Understanding; Shape-from-X; Probabilistic Model

1 Introduction

Understanding the structure of real-world 3D environments is fundamental to many computer vision tasks, with a well-studied history of research into 3D reconstruction from various visual input mediums, such as RGB video [43,13,55,53], RGB-D video [45,6,47,66,9], or single images [27,46,50,72,14,35,36,8]. Such approaches with active cameras have shown impressive capture of geometric structures leveraging strong visual signals. We consider an unconventional view of passive 3D scene perception: in the case of a lack of any visual signal, we look to human pose data, which for instance can be estimated from wearable IMU sensors [62,19,28], and ask “What can we learn about a 3D environment from only human pose trajectory information?”. This opens up new possibilities to explore information embedded in wearable devices (e.g., phones, fitness watches, etc.) towards understanding mapping, interactions, and content creation.




Fig. 1: From an observed pose trajectory of a person performing daily activities in an indoor scene (left), we learn to estimate likely object configurations of the scene underlying these interactions, as set of object class labels and oriented 3D bounding boxes (middle). By sampling from our probabilistic decoder, we synthesize multiple plausible object arrangements (right). (Scene geometry is shown only for visualization.)

In particular, we observe that human movement in a 3D environment often interacts both passively and actively with objects in the environment, giving strong cues about likely objects and their locations. For instance, walking around a room indicates where empty floor space is available, a sitting motion indicates high likelihood of a chair or sofa to support the sitting pose, and a single outstretched arm suggests picking up/putting down an object to furniture that supports the object. We thus propose to address a new scene estimation task: from only a sequence observation of 3D human poses, to estimate the object arrangement in the scene of the objects the person has interacted with, as a set of object class categories and 3D oriented bounding boxes (see Fig. 1).

As there are inherent ambiguities that lie in 3D object localization from only a human pose trajectory in the scene, we propose P2R-Net to learn a probabilistic model of the most likely modes of object configurations in the scene. From the sequence of poses, P2R-Net leverages the pose joint locations to vote for potential object centers that participate in the observed pose interactions. We then introduce a probabilistic decoder that learns a Gaussian mixture model for object box parameters, from which we can sample multiple diverse hypotheses of object arrangements. To enable massive training, we introduce a large-scale dataset with VirtualHome [51] platform to learn object configurations from human motions. Experiments on VirtualHome and the real dataset PROX [24,70] demonstrate our superiority against the baseline methods.

In summary, we present the following contributions:

- We propose a new perspective on 3D scene understanding by studying estimation of 3D object configurations from solely observing 3D human pose sequences of interactions in an environment, without any visual input, and predicting the object class categories and 3D oriented bounding boxes of the interacted objects in the scene.
- To address this task, we introduce a new, end-to-end, learned probabilistic model that estimates probability distributions for the object class categories and bounding box parameters.

- We demonstrate that our model captures complex, multi-modal distributions of likely object configurations, which can be sampled to produce diverse hypotheses that have accurate coverage over the ground truth object arrangement. Experiments also demonstrate the superiority of our method against the baselines in terms of accuracy and diversity.

2 Related Work

Predicting Human Interactions in Scenes. Capturing and modeling interactions between human and scenes has seen impressive progress in recent years, following significant advances in 3D reconstruction and 3D deep learning. From a visual observation of a scene, interactions and human-object relations are estimated. Several methods have been proposed for understanding the relations between scene and human poses via object functionality prediction [20,75,49,26] and affordance analysis [21,57,58,64,12,54].

By parsing the physics and semantics in human interactions, further works have been proposed towards synthesizing static human poses or human body models into 3D scenes [20,21,56,34,57,24,73,74,25]. These works focus on how to place human avatars into a 3D scene with semantic and physical plausibility (e.g., support or occlusion constraints). Various approaches have additionally explored synthesizing dynamic motions for a given scene geometry. Early methods retrieve and integrate existing avatar motions from database to make them compatible with scene geometry [37,2,32,38,60]. Given a goal pose or a task, more works learn to search for a possible motion path and estimate plausible contact motions [61,5,7,41,63,23]. These methods explore human-scene interaction understanding by estimating object functionalities or human interactions as poses in a given 3D scene environment. In contrast, we take a converse perspective, and aim to estimate the 3D scene arrangement from human pose trajectory observations.

Scene Understanding with Human Priors. As many environments, particularly indoor scenes, have been designed for people’s daily usage, human behavioral priors can be leveraged to additionally reason about 2D or 3D scene observations. Various methods have been proposed to leverage human context as extra signal towards holistic perception to improve performance in scene understanding tasks such as semantic segmentation [10], layout detection from images [16,59], 3D object labeling [29], 3D object detection and segmentation [65], and 3D reconstruction [17,18,70].

Additionally, several methods learn joint distributions of human interactions with 3D scenes or RGB video that can be leveraged to re-synthesize the observed scene as an arranged set of synthetic, labeled CAD models [31,15,30,57,42]. Recently, HPS [22] proposed to simultaneously estimate pose trajectory and scene reconstruction from wearable visual and inertial sensors on a person. We also aim to understand 3D scenes as arrangements of objects, but do not require any labeled interactions nor consider any visual (RGB, RGB-D, etc.) information as input. The recent approach of Mura et al. [44] poses the task of floor plan estimation from 2D human walk trajectories, and proposes to predict occupancy-based

floor plans that indicate structure and object footprints, but without object instance distinction and employs a fully-deterministic prediction. To the best of our knowledge, we introduce the first method to learn 3D object arrangement distributions from only human pose trajectories, without any visual input.

Pose Tracking with IMUs. Our method takes the input of human pose trajectories, which is built on the success of motion tracking techniques. Seminal work on pose estimation from wearable sensors have demonstrated effective pose estimation from wearable sensors, such as optical markers [4,23] or IMUs [39,62,28,33,19]. Our work is motivated by the capability of reliably estimating human pose from these sensor setups without visual data, from which we aim to learn human-object interaction priors to estimate scene object configurations.

3 Method

From only a human pose trajectory as input, we aim to estimate a distribution of likely object configurations, from which we can sample plausible hypotheses of objects in the scene as sets of class category labels and oriented 3D bounding boxes. We observe that most human interactions in an environment are targeted towards specific objects, and that general motion behavior is often influenced by the object arrangement in the scene. We thus aim to discover potential objects that each pose may be interacting with.

We first extract meaningful features from the human pose sequence with a *position encoder* to disentangle each frame into a relative position encoding and a position-agnostic pose, as well as a *pose encoder* to learn the local spatio-temporal feature for each pose in consecutive frames. We then leverage these features to vote for a potential interacting object for each pose. From these votes, we learn a *probabilistic mixture decoder* to propose box proposals for each object, characterizing likely modes for objectness, class label, and box parameters. An illustration of our approach is shown in Fig. 2.

3.1 Relative Position Encoding

We consider an input pose trajectory with N frames and J joints as the sequence of 3D locations $\mathbf{T} \in \mathbb{R}^{N \times J \times 3}$. We also denote the root joint of each pose by $\mathbf{r} \in \mathbb{R}^{N \times 3}$, where the root joint of a pose is the centroid of the joints corresponding to the body hip (for the skeleton configuration, we refer to the supplemental). To learn informative pose features, we first disentangle for each frame the absolute pose joint coordinates into a relative position encoding $\mathbf{Q} \in \mathbb{R}^{N \times d_1}$ and a position-agnostic pose feature $\mathbf{P} \in \mathbb{R}^{N \times J \times d_1}$, which are formulated as:

$$\begin{aligned} \mathbf{Q} &= \text{Pool}[f_1(\mathcal{N}(\mathbf{r}) - \mathbf{r})], \\ \mathbf{P} &= f_2(\mathbf{T} - \mathbf{r}), \quad \mathcal{N}(\mathbf{r}) \in \mathbb{R}^{N \times k \times 3}, \end{aligned} \tag{1}$$

where $f_1(*)$, $f_2(*)$ are point-wise MLP layers. $\mathcal{N}(\mathbf{r})$ is the set of k temporal neighbors to each root joint in \mathbf{r} , and $\text{Pool}(*)$ denotes neighbor-wise average

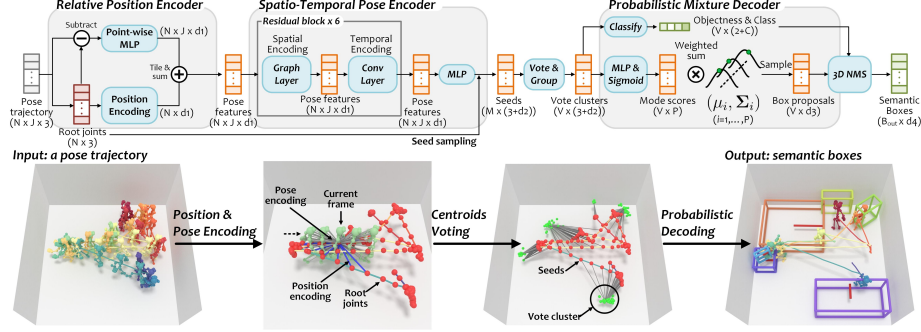


Fig. 2: Overview of P2R-Net. Given a pose trajectory with N frames and J joints, a position encoder decouples each skeleton frame into a relative position encoding (from its root joint as the hip centroid) and a position-agnostic pose. After combining them, a pose encoder learns local pose features from both body joints per skeleton (spatial encoding) and their changes in consecutive frames (temporal encoding). Root joints as seeds are then used to vote for the center of a nearby object that each pose is potentially interacting with. A probabilistic mixture network learns likely object box distributions, from which object class labels and oriented 3D boxes can be sampled.

pooling. By broadcast summation, we output $\mathbf{P}^r = \mathbf{P} + \mathbf{Q}$ for further spatio-temporal pose encoding. Understanding relative positions rather than absolute positions helps to provide more generalizable features to understand common pose motion in various object interactions, as these human-object interactions typically occur locally in a scene.

3.2 Spatio-Temporal Pose Encoding

The encoding \mathbf{P}^r provides signal for the relative pose trajectory of a person. We then further encode these features to capture the joint movement to understand local human-object interactions. That is, from \mathbf{P}^r , we learn joint movement in spatio-temporal domain: (1) in the spatial domain, we learn from intra-skeleton joints to capture per-frame pose features; (2) in the temporal domain, we learn from inter-frame relations to perceive each joint’s movement.

Inspired by [69] in 2D pose recognition, we first use a graph convolution layer to learn intra-skeleton joint features. Edges in the graph convolution are constructed following the skeleton bones, which encodes skeleton-wise spatial information. For each joint, we then use a 1-D convolution layer to capture temporal features from its inter-frame neighbors. A graph layer and an 1-D convolution layer are linked into a block with a residual connection to process the input \mathbf{P}^r (see Fig. 3). By stacking six blocks, we obtain a deeper spatio-temporal pose encoder with a wider receptive field in temporal domain, enabling reasoning over more temporal neighbors for object box estimation. Finally, we adopt an MLP to process all joints per skeleton to obtain pose features $\mathbf{P}^{st} \in \mathbb{R}^{N \times d_2}$.




Fig. 3: Pose encoding with spatio-temporal convolutions.

3.3 Locality-Sensitive Voting

With pose features \mathbf{P}^{st} , we then learn to vote for all the objects a person could have interacted with in a trajectory (see Fig. 2). For each pose frame, we predict the center of an object it potentially interacts with. Since we do not know when interactions begin or end, each pose votes for a potential object interaction. As human motion in a scene will tend to active interaction with an object or movement to an object, we aim to learn these patterns by encouraging votes for objects close to the respective pose, encouraging locality-based consideration.

For each pose feature $p^{st} \in \mathbf{P}^{st}$, we use its root joint $r \in \mathbf{r}$ as a seed location, and vote for an object center by learning the displacement from the seed:

$$\begin{aligned} \mathbf{v} &= \mathbf{r}_s + f_3(\mathbf{P}_s^{st}), \quad \mathbf{r}_s, \mathbf{v} \in \mathbb{R}^{M \times 3}, \\ \mathbf{P}^v &= \mathbf{P}_s^{st} + f_4(\mathbf{P}_s^{st}), \quad \mathbf{P}_s^{st}, \mathbf{P}^v \in \mathbb{R}^{M \times d_2}, \end{aligned} \quad (2)$$

where \mathbf{r}_s are the evenly sampled M seeds from \mathbf{r} ; \mathbf{P}_s^{st} are the corresponding pose features of \mathbf{r}_s ; f_3, f_4 are MLP layers; \mathbf{v}, \mathbf{P}^v denote the vote coordinates and features learned from \mathbf{P}_s^{st} . We evenly sample seeds \mathbf{r}_s from \mathbf{r} to make them cover the whole trajectory and adaptive to sequences with different length.

Since there are several objects in a scene, for each seed in \mathbf{r}_s , we vote for the center to the nearest one (see Fig. 4). The nearest object is both likely to participate in a nearby interaction, and affect motion around the object if not directly participating in an interaction. This strategy helps to reduce the ambiguities in scene object configuration estimation from a pose trajectory by capturing both direct and indirect effects of object location on pose movement.

For the seeds which vote for the same object, we group their votes to a cluster following [52]. This outputs cluster centers $\mathbf{v}^c \in \mathbb{R}^{V \times 3}$ with aggregated cluster feature $\mathbf{P}^c \in \mathbb{R}^{V \times d_2}$ where V denotes the number of vote clusters. We then use the \mathbf{P}^c to decode to distributions that characterize semantic 3D boxes, which is described in Section 3.4. For poses whose root joint is not close to any object during training (beyond a distance threshold t_d), we consider them to have little connection with the objects, and do not train them to vote for any object.

3.4 Probabilistic Mixture Decoder

We decode vote clusters $(\mathbf{v}^c, \mathbf{P}^c)$ to propose oriented 3D bounding boxes for each object, along with their class label and objectness score. Each box is represented




Fig. 4: Voting to objects that potentially influence the motion trajectory in approaching the target.

by a 3D center c , 3D size s and 1D orientation θ , where we represent the size by $\log(s)$ and orientation by $(\sin(\theta), \cos(\theta))$ for regression, similar to [71]. Since the nature of our task is inherently ambiguous (e.g., it may be unclear from observing a person sit if they are sitting on a chair or a sofa, or the size of the sofa), we propose to learn a probabilistic mixture decoder to predict the box centroid, size and orientation with multiple modes, from a vote cluster $v^c \in \mathbf{v}^c, P^c \in \mathbf{P}^c$:

$$\begin{aligned} y_\tau &= \sum_{k=1}^P f_\tau^k(P^c) \cdot y_\tau^k, \quad \tau \in \{c, s, \theta\}, \\ y_c^* &= v^c + y_c, \quad y_\tau^k \sim \mathcal{N}(\mu_\tau^k, \Sigma_\tau^k), \quad y_\tau, y_\tau^k \in \mathbb{R}^{d_\tau}, \end{aligned} \quad (3)$$

where $\tau \in \{c, s, \theta\}$ denote the regression targets for center, size, and orientation; $\mathcal{N}(\mu_\tau^k, \Sigma_\tau^k)$ is the learned multivariate Gaussian distribution of the k -th mode for τ , where y_τ^k is sampled from; P is the number of Gaussian distributions (i.e., modes); $f_\tau^k(*) \in [0, 1]$ is the learned score for the k -th mode; y_τ is the weighted sum of the samples from all modes, which is the prediction of the center/size/orientation; and d_τ is their output dimension ($d_c=3, d_s=3, d_\theta=2$). Note that the box center y_c^* is obtained by regressing the offset y_c from cluster center v^c . We predict the proposal objectness and the probability distribution for class category directly from P^c , using an MLP.

Multi-modal Prediction. In Eq. 3, the learnable parameters are $f_\tau(*)$ and (μ_τ, Σ_τ) . $f_\tau(*)$ is realized with an MLP followed by a sigmoid function, and (μ_τ, Σ_τ) are the learned embeddings shared among all samples. During training, we sample y_τ^k from each mode $\mathcal{N}(\mu_\tau^k, \Sigma_\tau^k)$ and predict y_τ using Eq. 3. To generate diverse and plausible hypotheses during inference, we not only sample y_τ^k , but also sample various different modes by randomly disregarding mixture elements based on their probabilities $f_\tau(*)$. Then we obtain y_τ as follows:

$$y_\tau = \sum_{k=1}^P I_\tau^k \cdot y_\tau^k, \quad I_\tau^k \sim \text{Bern}(f_\tau^k), \quad y_\tau^k \sim \mathcal{N}(\mu_\tau^k, \Sigma_\tau^k), \quad (4)$$

where I_τ^k is sampled from Bernoulli distribution with probability of $f_\tau^k(*)$. We also sample the object classes by the predicted classification probabilities, and discard proposed object boxes with low objectness ($\leq t_o$) after 3D NMS.

We can then generate N_h hypotheses in a scene; each hypothesis is an average of N_s samples of y_τ , which empirically strikes a good balance between diversity and accuracy of the set of hypotheses. To obtain the maximum likelihood pre-

diction, we use $f_{\tau}^k(*)$ and the mean value μ_{τ}^k instead of I_{τ}^k and y_{τ}^k to estimate the boxes with Eq. 3.

3.5 Loss Function

The loss consists of classification losses for objectness \mathcal{L}_{obj} and class label \mathcal{L}_{cls} , and regression losses for votes \mathcal{L}_v , box center \mathcal{L}_c , size \mathcal{L}_s and orientation \mathcal{L}_{θ} .

Classification Losses. Similar to [52], \mathcal{L}_{obj} and \mathcal{L}_{cls} are supervised by cross entropy losses, wherein the objectness score is used to classify if a vote cluster center is close to (≤ 0.3 m, positive) or far from (≥ 0.6 m, negative) the ground truth. Proposals from the clusters with positive objectness are further supervised with box regression losses.

Regression Losses. We supervise all the predicted votes, box centers, sizes and orientations with a Huber loss. For poses that are located within d_p to objects ($d_p = 1$ m), we use the closest object center to supervise their vote. Votes from those poses that are far from all objects are not considered. For center predictions, we use their nearest ground-truth center to calculate \mathcal{L}_c . Since box sizes and orientations are predicted from vote clusters, we use the counterpart from the ground-truth box that is nearest to the vote cluster for supervision. Then the final loss function is $\mathcal{L} = \sum_{\tau} \lambda_{\tau} \mathcal{L}_{\tau}$, where $\mathcal{L}_{\tau} \in \{\mathcal{L}_{obj}, \mathcal{L}_{cls}, \mathcal{L}_v, \mathcal{L}_c, \mathcal{L}_s, \mathcal{L}_{\theta}\}$ and $\{\lambda_{\tau}\}$ are constant weights that balance the losses.

4 Experiment Setup

Datasets. To the best of our knowledge, existing 3D human pose trajectory datasets are either with very few sequences (≤ 102) [57,24,74,3,70], or without instance annotations [24,3,74], or focused on single objects [23]. To this end, we introduce a new large-scale dataset using the simulation platform **VirtualHome** [51] for the task of estimating multiple scene objects from a human pose trajectory observation. To demonstrate our applicability to real data, we also evaluate on real human motions from the **PROX** dataset [24,70].

We construct our dataset on VirtualHome [51], which is built on the Unity3D game engine. It consists of 29 rooms, with each room containing 88 objects on average; each object is annotated with available interaction types. VirtualHome allows customization of action scripts and control over humanoid agents to execute a series of complex interactive tasks. We refer readers to [51] for the details of the scene and action types. In our work, we focus on the interactable objects under 17 common class categories. In each room, we select up to 10 random static objects to define the scene, and script the agent to interact with each of the objects in a sequential fashion. For each object, we also select a random interaction type associated with the object class category. Then we randomly sample 13,913 different sequences with corresponding object boxes to construct the dataset. During training, we also randomly flip, rotate and translate the scenes and poses for data augmentation. For additional detail about data generation, we refer to the supplemental. For the PROX dataset [24], we use its

human motions with the 3D instance boxes labeled by [70], which has 46 motion sequences interacting with four object categories (i.e., bed, chair, sofa, table) in 10 rooms. For more details, we refer readers to [24,70].

Implementation. We train P2R-Net end-to-end from scratch with the batch size at 32 on 4 NVIDIA 2080 Ti GPUs for 180 epochs, where Adam is used as the optimizer. The initial learning rate is at 1e-3 in the first 80 epochs, which is decayed by $0.1\times$ every 40 epochs after that. The losses are weighted by $\lambda_{obj}=5$, $\lambda_{cls}=1$, $\{\lambda_v, \lambda_c, \lambda_s, \lambda_\theta\}=10$ to balance the loss values. During training, we use pose distance threshold $t_d=1$ m. At inference time, we output box predictions after 3D NMS with an IoU threshold of 0.1. We use an objectness threshold of $t_o=0.5$. For the layer and data specifications, we refer to the supplemental.

Evaluation. We evaluate our task both on our dataset and on PROX. For our dataset, we consider two types of evaluation splits: a sequence-level split \mathcal{S}_1 across different interaction sequences, and room-level split \mathcal{S}_2 across different rooms as well as interaction sequences. Note that sequences are trained with and evaluated against only the objects that are interacted with during the input observation, resulting in different variants of each room under different interaction sequences. For \mathcal{S}_1 , the train/test split ratio is 4:1 over the generated sequences. \mathcal{S}_2 is a more challenging setup, with 27 train rooms and 2 test rooms, resulting in 13K/1K sequences. Since the task is inherently ambiguous, and only a single ground truth configuration of each room is available, we evaluate multi-modal predictions by several metrics: $mAP@0.5$ evaluates the mean average precision with the IoU threshold at 0.5 of the maximum likelihood prediction; MMD evaluates the Minimal Matching Distance [1] of the best matching prediction with the ground truth out of 10 sampled hypotheses to measure their quality; TMD evaluates the Total Mutual Diversity [68] to measure the diversity of the 10 hypotheses. We provide additional detail about MMD and TMD in the supplemental. For PROX dataset, we split the train/test set by 8:1 considering the very limited number of sequences (46) and use $mAP@0.5$ for detection evaluation.

5 Results and Analysis

We evaluate our approach on the task of scene object configuration estimation from a pose trajectory observation, in comparison with baselines constructed from state-of-the-art 3D detection and pose understanding methods, as well as an ablation analysis of our multi-modal prediction.

5.1 Baselines

Since there are no prior works that tackle the task of predicting the object configuration of a scene from solely a 3D pose trajectory, we construct several baselines leveraging state-of-the-art techniques as well as various approaches to estimate multi-modal distributions. We consider the following baselines: **1) Pose-VoteNet** [52]. Since VoteNet is designed for detection from point clouds, we replace their PointNet++ encoder with our position encoder + MLPs to

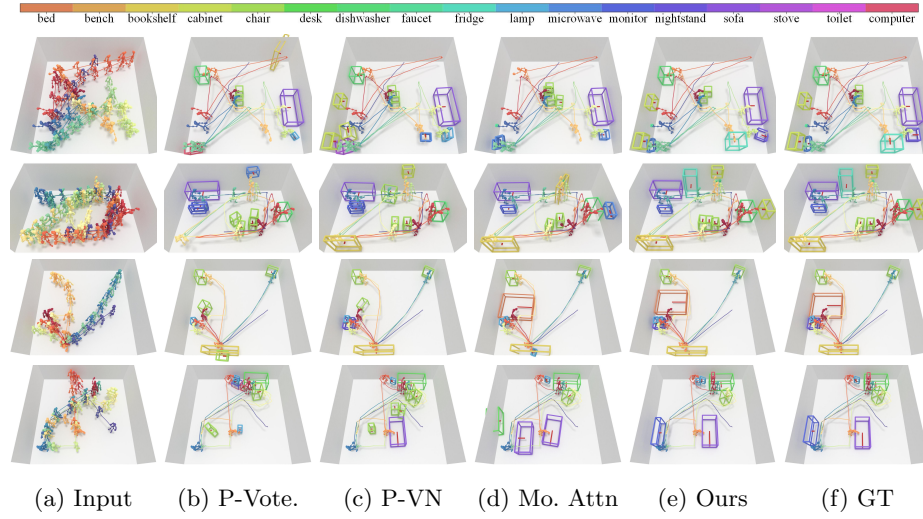


Fig. 5: Qualitative results of object detection from a pose trajectory on the sequence-level split \mathcal{S}_1 (unseen interaction sequences).

learn joint features for seeds. **2) Pose-VN**, Pose-VoteNet based on Vector Neurons [11] which replaces MLP layers in Pose-VoteNet with $SO(3)$ -equivariant operators that can capture arbitrary rotations of poses to estimate objects. **3) Motion Attention** [40]. Since our task can be also regarded as a sequence-to-sequence problem, we adopt a frame-wise attention encoder to extract repetitive pose patterns in the temporal domain which then inform a VoteNet decoder to regress boxes. Additionally, we also ablate our probabilistic mixture decoder with other alternatives: **4) Deterministic P2R-Net** (P2R-Net-D), where we use VoteNet decoder [52] in our method for box regression to produce deterministic results; **5) Generative P2R-Net** (P2R-Net-G), where our P2R-Net decoder is designed with a probabilistic generative model [67] to decode boxes from a learned latent variable; **6) Heatmap P2R-Net** (P2R-Net-H), where the box center, size and orientation are discretized into binary heatmaps, and the box regression is converted into a classification task. Detailed architecture specifications for these networks are given in the supplemental material.

5.2 Qualitative Comparisons

Comparisons on \mathcal{S}_1 . Fig. 5 visualizes predictions on the test set of unseen interaction sequences. Pose-VoteNet struggles to identify the existence of an object, leading to many missed detections, but can estimate reasonable object locations when an object is predicted. Pose-VN alleviates this problem of under-detection, but struggles to estimate object box sizes (rows 1,3). These baselines indicate the difficulty in detecting objects without sharing pose features among temporal neighbors. Motion Attention [40] addresses this by involving global context

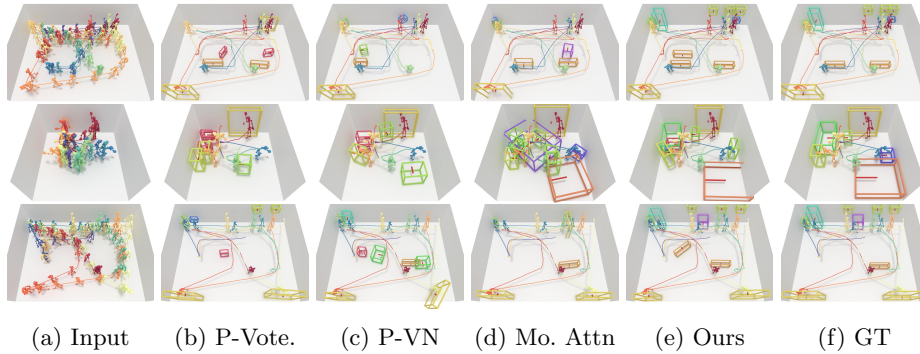


Fig. 6: Qualitative results of object detection from a pose trajectory on the room-level split \mathcal{S}_2 (unseen interaction sequences and rooms).

with inter-frame attention. However, it does not take advantage of the skeleton’s spatial layout in continuous frames and struggles to detect the existence of objects (row 1,2). In contrast, our method leverages both target-dependent poses and object occupancy context, that learns the implicit interrelations between poses and objects to infer object boxes, and achieves better estimate of the scene configuration.

Comparisons on \mathcal{S}_2 . In Fig. 6, we illustrate the qualitative comparisons on the test set of unseen interaction sequences in unknown rooms. In this scenario, most baselines fail to localize objects, while our method can nonetheless produce plausible object layouts.

Multi-modal Predictions. We visualize various sampled hypotheses from our model \mathcal{S}_1 in Fig. 8, showing that our method is able to deduce the spatial occupancy of objects from motion trajectories, and enables diverse, plausible estimation of object locations, orientation, and sizes for interactions.

Comparisons on PROX. We additionally show qualitative results on real motion data from PROX [24,70] in Fig.7. Since there are only 46 sequences labeled with object boxes, we pretrain each method on our dataset first before

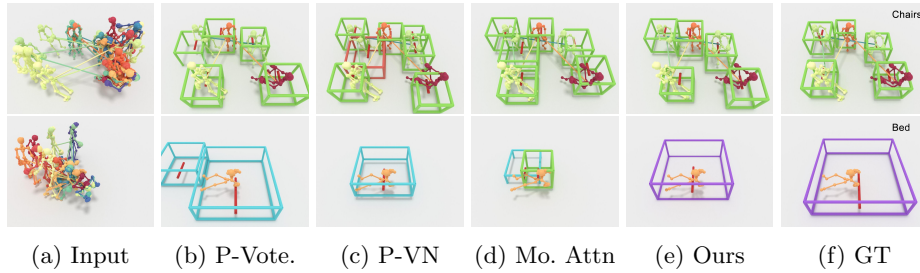


Fig. 7: Qualitative results of object detection on the PROX dataset [24,70].




Fig. 8: Multi-modal predictions of P2R-Net. By sampling our decoder multiple times, we can obtain different plausible box predictions. Here, we show three randomly sampled hypotheses and the max. likelihood prediction for each input.

| | bed | bench | cabinet | chair | desk | dishwasher | fridge | lamp | sofa | stove | toilet | computer | mAP@0.5 |
|------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|-------------|--------------|--------------|--------------|--------------|--------------|
| Pose-VoteNet | 2.90 | 15.00 | 33.14 | 18.77 | 58.52 | 32.14 | 0.00 | 6.07 | 62.32 | 49.82 | 0.00 | 3.06 | 25.70 |
| Pose-VN | 20.81 | 18.13 | 49.76 | 18.68 | 70.92 | 33.56 | 0.00 | 5.60 | 67.24 | 46.76 | 0.00 | 6.11 | 29.90 |
| Motion Attention | 36.42 | 7.54 | 23.35 | 19.50 | 77.71 | 15.59 | 17.13 | 2.35 | 78.61 | 51.03 | 14.81 | 5.50 | 28.39 |
| P2R-Net-D | 93.77 | 12.63 | 11.98 | 5.77 | 95.93 | 61.80 | 73.95 | 0.58 | 88.44 | 70.42 | 0.00 | 14.17 | 34.91 |
| P2R-Net-G | 91.69 | 7.56 | 36.61 | 10.05 | 93.47 | 67.53 | 77.45 | 1.21 | 92.97 | 64.86 | 5.56 | 18.67 | 37.48 |
| P2R-Net-H | 85.84 | 8.04 | 22.04 | 10.91 | 76.08 | 55.20 | 55.15 | 0.00 | 83.92 | 57.60 | 5.00 | 4.33 | 31.41 |
| P2R-Net | 94.21 | 10.12 | 54.72 | 8.02 | 93.32 | 56.33 | 59.89 | 3.25 | 90.92 | 57.86 | 61.11 | 19.94 | 42.20 |

Table 1: Quantitative evaluation on split \mathcal{S}_1 . For P2R-Net-G, P2R-Net-H and ours, we use the maximum likelihood predictions to calculate mAP scores. The mAP@0.5 is averaged over all 17 classes (see the full table with all classes in the supplementary file).

finetuning them on PROX. As PROX uses SMPL-X human body model [48], we manually map its skeleton joint indexes to ours for transfer learning. The results show that our approach can effectively handle real, noisy pose trajectory inputs.

5.3 Quantitative Comparisons

We use mAP@0.5 to measure object detection accuracy, and evaluate the accuracy and diversity of a set of output hypotheses using minimal matching distance (MMD) and total mutual diversity (TMD).

Detection Accuracy. Table 1 shows a quantitative comparison on split \mathcal{S}_1 , where we observe that Pose-VoteNet and Pose-VN, struggle to recognize some object categories (e.g., bed, fridge and toilet). By leveraging the inter-frame connections, Motion Attention achieves improved performance in recognizing object classes, but struggles with detecting objectness and predicting object sizes. In contrast, our position and pose encoder learns both the spatial and temporal signals from motions to estimate likely object locations by leveraging the potential connections between human and objects, with our probabilistic mixture decoder better capturing likely modalities in various challenging ambiguous interactions (e.g., toilet and cabinet). In Table 2, we compare mAP@0.5 scores on split \mathcal{S}_2 , with increased relative improvement in the challenging scenario of scene object configuration estimation in new rooms.

| | \mathcal{S}_1 | | | \mathcal{S}_2 | | |
|--------------|-----------------|--------------|-------------|-----------------|--------------|-------------|
| | mAP | MMD | TMD | mAP | MMD | TMD |
| Pose-VoteNet | 25.70 | - | - | 10.23 | - | - |
| Pose-VN | 29.90 | - | - | 6.12 | - | - |
| Motion Attn. | 28.39 | - | - | 6.53 | - | - |
| P2R-Net-D | 34.91 | - | - | 31.56 | - | - |
| P2R-Net-G | 37.48 | 37.43 | 1.73 | 31.59 | 31.83 | 1.98 |
| P2R-Net-H | 31.41 | 21.10 | 2.39 | 27.96 | 20.71 | 3.28 |
| P2R-Net | 42.20 | 38.28 | 3.34 | 35.34 | 32.47 | 4.35 |

Table 2: Comparisons on detection accuracy, and multi-modal quality and diversity on \mathcal{S}_1 and \mathcal{S}_2 . TMD=1 indicates no diversity.

| | w/o pretrain | w/ pretrain |
|--------------|--------------|--------------|
| Pose-VoteNet | 2.29 | 27.97 |
| Pose-VN | 1.92 | 9.23 |
| Motion Attn. | 2.68 | 6.25 |
| P2R-Net | 5.36 | 31.38 |

Table 3: Comparisons on detection accuracy (mAP@0.5) using PROX.

Quality and Diversity of Multi-modal Predictions. We study the multi-modal predictions with our ablation variants P2R-Net-G and P2R-Net-H, and use MMD and TMD to evaluate the quality and diversity of 10 randomly sampled predictions for each method. From the 10 predictions, MMD records the best detection score (mAP@0.5), and TMD measures the average variance of the 10 semantic boxes per object. Table 2 presents the MMD and TMD scores on \mathcal{S}_1 and \mathcal{S}_2 respectively. P2R-Net-G tends to predict with low diversity, as seen in low TMD score (TMD=1 indicates identical samples) and similar mAP@0.5 and MMD scores. P2R-Net-H shows better diversity but with lower accuracy in both of the two splits. Our final model not only achieves best detection accuracy, it also provides reasonable and diverse object configurations, with a notable performance improvement in the challenging \mathcal{S}_2 split.

Comparisons on PROX. We quantitatively compare with baselines on real human motion data from PROX [24,70] in Table 3. We evaluate our method and baselines with and without pretraining on our synthetic dataset, considering the very limited number (46) of motion sequences in PROX. The results demonstrate that pretraining on our dataset can significantly improve all methods’ performance on real data, with our approach outperforming all baselines.

Ablations. In Table 4, we explore the effects of each individual module (relative position encoder, spatio-temporal pose encoder and probabilistic mixture decoder). We ablate P2R-Net by gradually adding components from the baseline \mathbf{c}_0 : without relative position encoding (Pstn-Enc), where we use joints’ global coordinates relative to the room center; without spatio-temporal pose encoding (Pose-Enc), where we use MLPs to learn pose features from joint coordinates; without probabilistic mixture decoder (P-Dec), where we use two-layer MLPs to regress box parameters.

From the comparisons, we observe that relative position encoding plays the most significant role. It allows our model to pick up on local pose signal, as many human-object interactions present with strong locality. The spatio-temporal pose encoder then enhances the pose feature learning, and enables our model to learn the joint changes both in spatial and temporal domains. This design largely improves our generalization ability, particularly in unseen rooms (from 19.07 to

| $\mathcal{S}_1/\mathcal{S}_2$ | Pstn-Enc | Pose-Enc | P-Dec | mAP@0.5 |
|-------------------------------|----------|----------|-------|----------------------|
| \mathbf{c}_0 | - | - | - | 8.71 / 3.12 |
| \mathbf{c}_1 | ✓ | - | - | 34.43 / 19.07 |
| \mathbf{c}_2 | ✓ | ✓ | - | 39.98 / 29.25 |
| Full | ✓ | ✓ | ✓ | 42.20 / 35.34 |

Table 4: Ablation study of our design choices. Note that \mathbf{c}_1 and \mathbf{c}_2 are different with Pose-VoteNet and P2R-Net-D since our method parameterizes boxes differently from VoteNet (see Sec. 3.4)

29.25). The probabilistic decoder further alleviates the ambiguity of this problem, and combining all the modules together produces the best results.

Limitations. Although P2R-Net achieves plausible scene estimations from only pose trajectories, it operates on several assumptions that can lead to potential limitations: (1) Objects should be interactable, e.g., our method may not detect objects that do not induce strong pose interactions like mirror or picture; (2) Interactions occur at close range, e.g., we may struggle to detect a TV from a person switching on it with a remote control. Additionally, we currently focus on estimating static object boxes. We believe an interesting avenue for future work is to characterize objects in motion (e.g., due to grabbing) or articulated objects (e.g., laptops).

6 Conclusion

We have presented a first exploration to the ill-posed problem of estimating the 3D object configuration in a scene from only a 3D pose trajectory observation of a person interacting with the scene. Our proposed model P2R-Net leverages spatio-temporal features from the pose trajectory to vote for likely object positions and inform a new probabilistic mixture decoder that captures multi-modal distributions of object box parameters. We demonstrate that such a probabilistic approach can effectively model likely object configurations in scene, producing plausible object layout hypotheses from an input pose trajectory. We hope that this establishes a step towards object-based 3D understanding of environments using non-visual signal and opens up new possibilities in leveraging ego-centric motion for 3D perception and understanding.

Acknowledgments

This project is funded by the Bavarian State Ministry of Science and the Arts and coordinated by the Bavarian Research Institute for Digital Transformation (bidt), the TUM Institute of Advanced Studies (TUM-IAS), the ERC Starting Grant Scan2CAD (804724), and the German Research Foundation (DFG) Grant Making Machine Learning on Static and Dynamic 3D Data Practical.

References

1. Achlioptas, P., Diamanti, O., Mitliagkas, I., Guibas, L.: Learning representations and generative models for 3d point clouds. In: International conference on machine learning. pp. 40–49. PMLR (2018) [9](#)
2. Agrawal, S., van de Panne, M.: Task-based locomotion. *ACM Transactions on Graphics (TOG)* **35**(4), 1–11 (2016) [3](#)
3. Cao, Z., Gao, H., Mangalam, K., Cai, Q.Z., Vo, M., Malik, J.: Long-term human motion prediction with scene context. In: European Conference on Computer Vision. pp. 387–404. Springer (2020) [8](#)
4. Chai, J., Hodgins, J.K.: Performance animation from low-dimensional control signals. In: ACM SIGGRAPH 2005 Papers, pp. 686–696 (2005) [4](#)
5. Chao, Y.W., Yang, J., Chen, W., Deng, J.: Learning to sit: Synthesizing human-chair interactions via hierarchical control. arXiv preprint arXiv:1908.07423 (2019) [3](#)
6. Choi, S., Zhou, Q.Y., Koltun, V.: Robust reconstruction of indoor scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5556–5565 (2015) [1](#)
7. Corona, E., Pumarola, A., Alenya, G., Moreno-Noguer, F.: Context-aware human motion prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6992–7001 (2020) [3](#)
8. Dahnert, M., Hou, J., Nießner, M., Dai, A.: Panoptic 3d scene reconstruction from a single rgb image. *Proc. Neural Information Processing Systems (NeurIPS)* (2021) [1](#)
9. Dai, A., Nießner, M., Zollhöfer, M., Izadi, S., Theobalt, C.: Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration. *ACM Transactions on Graphics (ToG)* **36**(4), 1 (2017) [1](#)
10. Delaitre, V., Fouhey, D.F., Laptev, I., Sivic, J., Gupta, A., Efros, A.A.: Scene semantics from long-term observation of people. In: European conference on computer vision. pp. 284–298. Springer (2012) [3](#)
11. Deng, C., Litany, O., Duan, Y., Poulenard, A., Tagliasacchi, A., Guibas, L.: Vector neurons: a general framework for so(3)-equivariant networks. arXiv preprint arXiv:2104.12229 (2021) [10](#)
12. Deng, S., Xu, X., Wu, C., Chen, K., Jia, K.: 3d affordancenet: A benchmark for visual object affordance understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1778–1787 (2021) [3](#)
13. Engel, J., Schöps, T., Cremers, D.: Lsd-slam: Large-scale direct monocular slam. In: European conference on computer vision. pp. 834–849. Springer (2014) [1](#)
14. Engelmann, F., Rematas, K., Leibe, B., Ferrari, V.: From points to multi-object 3d reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4588–4597 (2021) [1](#)
15. Fisher, M., Savva, M., Li, Y., Hanrahan, P., Nießner, M.: Activity-centric scene synthesis for functional 3d scene modeling. *ACM Transactions on Graphics (TOG)* **34**(6), 1–13 (2015) [3](#)
16. Fouhey, D.F., Delaitre, V., Gupta, A., Efros, A.A., Laptev, I., Sivic, J.: People watching: Human actions as a cue for single view geometry. In: European Conference on Computer Vision. pp. 732–745. Springer (2012) [3](#)
17. Fowler, S., Kim, H., Hilton, A.: Towards complete scene reconstruction from single-view depth and human motion. In: BMVC (2017) [3](#)

18. Fowler, S., Kim, H., Hilton, A.: Human-centric scene understanding from single view 360 video. In: 2018 International Conference on 3D Vision (3DV). pp. 334–342. IEEE (2018) [3](#)
19. Glauser, O., Wu, S., Panozzo, D., Hilliges, O., Sorkine-Hornung, O.: Interactive hand pose estimation using a stretch-sensing soft glove. *ACM Transactions on Graphics (TOG)* **38**(4), 1–15 (2019) [1](#), [4](#)
20. Grabner, H., Gall, J., Van Gool, L.: What makes a chair a chair? In: CVPR 2011. pp. 1529–1536. IEEE (2011) [3](#)
21. Gupta, A., Satkin, S., Efros, A.A., Hebert, M.: From 3d scene geometry to human workspace. In: CVPR 2011. pp. 1961–1968. IEEE (2011) [3](#)
22. Guzov, V., Mir, A., Sattler, T., Pons-Moll, G.: Human poseitioning system (hps): 3d human pose estimation and self-localization in large scenes from body-mounted sensors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4318–4329 (2021) [3](#)
23. Hassan, M., Ceylan, D., Villegas, R., Saito, J., Yang, J., Zhou, Y., Black, M.J.: Stochastic scene-aware motion prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11374–11384 (2021) [3](#), [4](#), [8](#)
24. Hassan, M., Choutas, V., Tzionas, D., Black, M.J.: Resolving 3d human pose ambiguities with 3d scene constraints. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2282–2292 (2019) [2](#), [3](#), [8](#), [9](#), [11](#), [13](#)
25. Hassan, M., Ghosh, P., Tesch, J., Tzionas, D., Black, M.J.: Populating 3d scenes by learning human-scene interaction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14708–14718 (2021) [3](#)
26. Hu, R., van Kaick, O., Wu, B., Huang, H., Shamir, A., Zhang, H.: Learning how objects function via co-analysis of interactions. *ACM Transactions on Graphics (TOG)* **35**(4), 1–13 (2016) [3](#)
27. Huang, S., Qi, S., Zhu, Y., Xiao, Y., Xu, Y., Zhu, S.C.: Holistic 3d scene parsing and reconstruction from a single rgb image. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 187–203 (2018) [1](#)
28. Huang, Y., Kaufmann, M., Aksan, E., Black, M.J., Hilliges, O., Pons-Moll, G.: Deep inertial poser: Learning to reconstruct human pose from sparse inertial measurements in real time. *ACM Transactions on Graphics (TOG)* **37**(6), 1–15 (2018) [1](#), [4](#)
29. Jiang, Y., Koppula, H., Saxena, A.: Hallucinated humans as the hidden context for labeling 3d scenes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2993–3000 (2013) [3](#)
30. Jiang, Y., Koppula, H.S., Saxena, A.: Modeling 3d environments through hidden human context. *IEEE transactions on pattern analysis and machine intelligence* **38**(10), 2040–2053 (2015) [3](#)
31. Jiang, Y., Lim, M., Saxena, A.: Learning object arrangements in 3d scenes using human context. *arXiv preprint arXiv:1206.6462* (2012) [3](#)
32. Kapadia, M., Xianghao, X., Nitti, M., Kallmann, M., Coros, S., Sumner, R.W., Gross, M.: Precision: Precomputing environment semantics for contact-rich character animation. In: Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games. pp. 29–37 (2016) [3](#)
33. Kaufmann, M., Zhao, Y., Tang, C., Tao, L., Twigg, C., Song, J., Wang, R., Hilliges, O.: Em-pose: 3d human pose estimation from sparse electromagnetic trackers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11510–11520 (2021) [4](#)
34. Kim, V.G., Chaudhuri, S., Guibas, L., Funkhouser, T.: Shape2pose: Human-centric shape analysis. *ACM Transactions on Graphics (TOG)* **33**(4), 1–12 (2014) [3](#)

35. Kuo, W., Angelova, A., Lin, T.Y., Dai, A.: Mask2cad: 3d shape prediction by learning to segment and retrieve. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16. pp. 260–277. Springer (2020) [1](#)
36. Kuo, W., Angelova, A., Lin, T.Y., Dai, A.: Patch2cad: Patchwise embedding learning for in-the-wild shape retrieval from a single image. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12589–12599 (2021) [1](#)
37. Lee, J., Chai, J., Reitsma, P.S., Hodgins, J.K., Pollard, N.S.: Interactive control of avatars animated with human motion data. In: Proceedings of the 29th annual conference on Computer graphics and interactive techniques. pp. 491–500 (2002) [3](#)
38. Lee, K.H., Choi, M.G., Lee, J.: Motion patches: building blocks for virtual environments annotated with motion data. In: ACM SIGGRAPH 2006 Papers, pp. 898–906 (2006) [3](#)
39. Liu, H., Wei, X., Chai, J., Ha, I., Rhee, T.: Realtime human motion control with a small number of inertial sensors. In: Symposium on interactive 3D graphics and games. pp. 133–140 (2011) [4](#)
40. Mao, W., Liu, M., Salzmann, M.: History repeats itself: Human motion prediction via motion attention. In: European Conference on Computer Vision. pp. 474–489. Springer (2020) [10](#)
41. Merel, J., Tunyasuvunakool, S., Ahuja, A., Tassa, Y., Hasenclever, L., Pham, V., Erez, T., Wayne, G., Heess, N.: Catch & carry: reusable neural controllers for vision-guided whole-body tasks. *ACM Transactions on Graphics (TOG)* **39**(4), 39–1 (2020) [3](#)
42. Monszpart, A., Guerrero, P., Ceylan, D., Yumer, E., Mitra, N.J.: imapper: interaction-guided scene mapping from monocular videos. *ACM Transactions On Graphics (TOG)* **38**(4), 1–15 (2019) [3](#)
43. Mur-Artal, R., Montiel, J.M.M., Tardos, J.D.: Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics* **31**(5), 1147–1163 (2015) [1](#)
44. Mura, C., Pajarola, R., Schindler, K., Mitra, N.: Walk2map: Extracting floor plans from indoor walk trajectories. In: Computer Graphics Forum. vol. 40, pp. 375–388. Wiley Online Library (2021) [3](#)
45. Newcombe, R.A., Izadi, S., Hilliges, O., Molyneaux, D., Kim, D., Davison, A.J., Kohi, P., Shotton, J., Hodges, S., Fitzgibbon, A.: Kinectfusion: Real-time dense surface mapping and tracking. In: 2011 10th IEEE international symposium on mixed and augmented reality. pp. 127–136. IEEE (2011) [1](#)
46. Nie, Y., Han, X., Guo, S., Zheng, Y., Chang, J., Zhang, J.J.: Total3dunderstanding: Joint layout, object pose and mesh reconstruction for indoor scenes from a single image. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020) [1](#)
47. Nießner, M., Zollhöfer, M., Izadi, S., Stamminger, M.: Real-time 3d reconstruction at scale using voxel hashing. *ACM Transactions on Graphics (ToG)* **32**(6), 1–11 (2013) [1](#)
48. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10975–10985 (2019) [12](#)
49. Pieropan, A., Ek, C.H., Kjellström, H.: Functional object descriptors for human activity modeling. In: 2013 IEEE International Conference on Robotics and Automation. pp. 1282–1289. IEEE (2013) [3](#)

50. Popov, S., Bauszat, P., Ferrari, V.: Corenet: Coherent 3d scene reconstruction from a single rgb image. In: European Conference on Computer Vision. pp. 366–383. Springer (2020) [1](#)
51. Puig, X., Ra, K., Boben, M., Li, J., Wang, T., Fidler, S., Torralba, A.: Virtual-home: Simulating household activities via programs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8494–8502 (2018) [2](#), [8](#)
52. Qi, C.R., Litany, O., He, K., Guibas, L.J.: Deep hough voting for 3d object detection in point clouds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2019) [6](#), [8](#), [9](#), [10](#)
53. Qian, S., Jin, L., Fouhey, D.F.: Associative3d: Volumetric reconstruction from sparse views. In: European Conference on Computer Vision. pp. 140–157. Springer (2020) [1](#)
54. Ruiz, E., Mayol-Cuevas, W.: Where can i do this? geometric affordances from a single example with the interaction tensor. In: 2018 IEEE International Conference on Robotics and Automation (ICRA). pp. 2192–2199. IEEE (2018) [3](#)
55. Runz, M., Li, K., Tang, M., Ma, L., Kong, C., Schmidt, T., Reid, I., Agapito, L., Straub, J., Lovegrove, S., et al.: Frodo: From detections to 3d objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14720–14729 (2020) [1](#)
56. Savva, M., Chang, A.X., Hanrahan, P., Fisher, M., Nießner, M.: Scenegrok: Inferring action maps in 3d environments. *ACM transactions on graphics (TOG)* **33**(6), 1–10 (2014) [3](#)
57. Savva, M., Chang, A.X., Hanrahan, P., Fisher, M., Nießner, M.: Pigraphs: learning interaction snapshots from observations. *ACM Transactions on Graphics (TOG)* **35**(4), 1–12 (2016) [3](#), [8](#)
58. Sawatzky, J., Srikantha, A., Gall, J.: Weakly supervised affordance detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2795–2804 (2017) [3](#)
59. Shoaib, M., Yang, M.Y., Rosenhahn, B., Ostermann, J.: Estimating layout of cluttered indoor scenes using trajectory-based priors. *Image and Vision Computing* **32**(11), 870–883 (2014) [3](#)
60. Shum, H.P., Komura, T., Shiraishi, M., Yamazaki, S.: Interaction patches for multi-character animation. *ACM Transactions on Graphics (TOG)* **27**(5), 1–8 (2008) [3](#)
61. Starke, S., Zhang, H., Komura, T., Saito, J.: Neural state machine for character-scene interactions. *ACM Trans. Graph.* **38**(6), 209–1 (2019) [3](#)
62. Von Marcard, T., Rosenhahn, B., Black, M.J., Pons-Moll, G.: Sparse inertial poser: Automatic 3d human pose estimation from sparse imus. In: *Computer Graphics Forum*. vol. 36, pp. 349–360. Wiley Online Library (2017) [1](#), [4](#)
63. Wang, J., Xu, H., Xu, J., Liu, S., Wang, X.: Synthesizing long-term 3d human motion and interaction in 3d scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9401–9411 (2021) [3](#)
64. Wang, Z., Chen, L., Rathore, S., Shin, D., Fowlkes, C.: Geometric pose affordance: 3d human pose with scene constraints. *arXiv preprint arXiv:1905.07718* (2019) [3](#)
65. Wei, P., Zhao, Y., Zheng, N., Zhu, S.C.: Modeling 4d human-object interactions for joint event segmentation, recognition, and object localization. *IEEE transactions on pattern analysis and machine intelligence* **39**(6), 1165–1179 (2016) [3](#)
66. Whelan, T., Leutenegger, S., Salas-Moreno, R., Glocker, B., Davison, A.: Elastic-fusion: Dense slam without a pose graph. *Robotics: Science and Systems* (2015) [1](#)

67. Wu, J., Zhang, C., Xue, T., Freeman, W.T., Tenenbaum, J.B.: Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In: Proceedings of the 30th International Conference on Neural Information Processing Systems. pp. 82–90 (2016) [10](#)
68. Wu, R., Chen, X., Zhuang, Y., Chen, B.: Multimodal shape completion via conditional generative adversarial networks. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16. pp. 281–296. Springer (2020) [9](#)
69. Yan, S., Xiong, Y., Lin, D.: Spatial temporal graph convolutional networks for skeleton-based action recognition. In: Thirty-second AAAI conference on artificial intelligence (2018) [5](#)
70. Yi, H., Huang, C.H.P., Tzionas, D., Kocabas, M., Hassan, M., Tang, S., Thies, J., Black, M.J.: Human-aware object placement for visual environment reconstruction. In: Computer Vision and Pattern Recognition (CVPR) (2022) [2](#), [3](#), [8](#), [9](#), [11](#), [13](#)
71. Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3d object detection and tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11784–11793 (2021) [7](#)
72. Zhang, C., Cui, Z., Zhang, Y., Zeng, B., Pollefeys, M., Liu, S.: Holistic 3d scene understanding from a single image with implicit representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8833–8842 (June 2021) [1](#)
73. Zhang, S., Zhang, Y., Ma, Q., Black, M.J., Tang, S.: Place: Proximity learning of articulation and contact in 3d environments. In: 2020 International Conference on 3D Vision (3DV). pp. 642–651. IEEE (2020) [3](#)
74. Zhang, Y., Hassan, M., Neumann, H., Black, M.J., Tang, S.: Generating 3d people in scenes without people. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6194–6204 (2020) [3](#), [8](#)
75. Zhu, Y., Zhao, Y., Chun Zhu, S.: Understanding tools: Task-oriented object modeling, learning and recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2855–2864 (2015) [3](#)