

FutureHuman3D: Forecasting Complex Long-Term 3D Human Behavior from Video Observations

Christian Diller
Technical University of Munich
christian.diller@tum.de

Thomas Funkhouser
Google
tfunkhouser@google.com

Angela Dai
Technical University of Munich
angela.dai@tum.de




Figure 1. We propose a novel generative approach to model long-term future human behavior by jointly forecasting a sequence of coarse action labels and their concrete realizations as 3D body poses. For broad applicability, our autoregressive method only requires weak supervision and past observations in the form of 2D RGB video data, together with a database of uncorrelated 3D human poses.

Abstract

We present a generative approach to forecast long-term future human behavior in 3D, requiring only weak supervision from readily available 2D human action data. This is a fundamental task enabling many downstream applications. The required ground-truth data is hard to capture in 3D (mocap suits, expensive setups) but easy to acquire in 2D (simple RGB cameras). Thus, we design our method to only require 2D RGB data while being able to generate 3D human motion sequences. We use a differentiable 2D projection scheme in an autoregressive manner for weak supervision, and an adversarial loss for 3D regularization. Our method predicts long and complex behavior sequences (e.g. cooking, assembly) consisting of multiple sub-actions. We tackle this in a semantically hierarchical manner, jointly predicting high-level coarse action labels together with their low-level fine-grained realizations as characteristic 3D human poses. We observe that these two action representations are coupled in nature, and joint prediction benefits both action and pose forecasting. Our experiments demonstrate the complementary nature of joint action and 3D pose prediction: our joint approach outperforms each task treated individually, enables robust longer-term sequence prediction, and outperforms alternative approaches to forecast actions and characteristic 3D poses.

1. Introduction

Predicting future human behavior is fundamental to machine intelligence, with many applications in content creation, robotics, mixed reality, and more. For instance, a monitoring system might issue early warnings of potentially dangerous behaviour, and a robotic assistant can use forecasting to place tools at the right place and time they will be needed in the future. Consider the specific scenario of an assembly line monitoring system deployed to issue early warnings of behavior that could be harmful in the near future: The system needs to have a long-term understanding of the future, enabling it to forecast multiple action steps ahead so that it can act in time before a harmful action occurs. However, simply understanding the next action steps on a high level is not sufficient: it must also reason about *where* the action will occur. Actions such as “grab a tool” are likely harmless when performed in a toolbox; they become dangerous when done next to an active table saw or moving robot arm. The monitoring system thus also needs to be able to reason about spatial relations in 3D – for both the location and body pose of involved humans.

To support these types of applications, we must address two tasks: 1) forecasting long-term action sequences, and 2) predicting future 3D human poses. Prior work has focused on each of these tasks separately: activity forecasting predicts

future action labels without considering the 3D poses [33, 35, 36, 49, 51, 71], while 3D pose forecasting focuses on fixed frame rate sequence prediction limited to single actions in short-term time frames without considering longer-term action sequences [31, 61, 62, 93, 96].

We propose that these two tasks are coupled in nature: predicting action labels with realized 3D poses helps to encourage richer feature learning and can materialize sub-category level differences in actions for predicting future activities, and grounding 3D poses with actions provides global structure for longer-term forecasting.

Leveraging this insight, we design a method that takes in a sequence of recent RGB image observations and their action labels, and jointly predicts a sequence of future 3D characteristic poses and action labels (Fig. 1). In our design, we had to address two significant research challenges: 1) forecasting 3D poses from 2D images without any paired 3D training data, and 2) forecasting long sequences of actions comprising several discrete action steps.

The first challenge arises from limited training data. It would be ideal to have a dataset with ground truth 3D pose and action annotations for complex sequences of actions. Unfortunately, no such dataset exists. There are RGB video datasets with tracked 3D poses for limited types of actions (e.g., walking or waving); and there are video datasets with action labels for complex sequences of actions (e.g., cooking or assembly). However, there is no single dataset that has both types of annotations, and capturing one would be difficult due to the challenges of setting up 3D pose trackers in settings where people typically perform complex sequences of actions (e.g., cooking in a kitchen). Instead, we have to learn to use 2D video observations for 3D pose and action label forecasting without paired data. We achieve this by weakly supervising our pose forecasting in 2D using readily available 2D action datasets [8, 72] and formulate an adversarial loss encouraging likely 3D characteristic poses with respect to a distribution learned from 3D pose datasets [42, 60, 82]. Crucially, this does not require any correspondence between the 2D video and 3D pose data.

The second challenge arises from the difficulties of predicting long sequences of discrete events. One option would be to train a model to output a multi-step sequence of actions and poses all at once – however, this is impossible given the exponential growth of multi-step sequences and the limited amount of available training data. Another option would be to train a model that predicts the next future poses and actions at fixed time points in the future (e.g., 1s in advance) and then recurrently make long-term predictions – however, this time-based forecasting approach produces sequences that tend to “drift” over the long-term, since the intermediate poses at fixed time steps are usually “in between” semantically meaningful actions and thus do not provide a distinctive input representation for the next prediction. To

address this issue, we train our autoregressive approach to iteratively generate the next discrete action label along with the *3D characteristic pose* for that action. A 3D characteristic pose [22] is the set of 3D joint positions corresponding to the most distinctive moment a semantic action is being performed (e.g., when a hand grasps an object, when two objects are first brought together, etc.). By training our method to produce these poses as intermediate outputs (and inputs to the next step), we are able to generate more semantically plausible forecasts over longer action sequences.

Our experiments with two RGB video datasets demonstrate that our approach for joint prediction of action behaviors and 3D poses outperforms state-of-the-art methods applied separately to each task. Additionally, we find that predicting actions and their 3D characteristic poses enables more robust autoregressive prediction for longer-term forecasting. Overall, our contributions are:

- The first method to learn forecasting of future 3D poses from datasets with only 2D RGB video and action label data (i.e., without any paired 3D data).
- The first method to forecast future 3D poses jointly with action labels from commonly available video input.
- The first method to forecast future characteristic 3D poses and action labels for long-term and complex behaviors.

2. Related Work

3D Human Pose Forecasting. Forecasting 3D human poses has been studied in many previous works and is commonly formulated as a 3D sequential motion prediction task, taking an input 3D sequence of poses and generating an output 3D sequence of poses. For short-term future prediction (up to ≈ 1 second), RNN-based approaches have achieved impressive performance [1, 16, 30, 37, 38, 44, 64, 69, 90]. As RNNs tend to struggle to capture longer-term dependencies with a fixed-size history, graph-based [17, 18, 20, 53, 55, 56, 61, 77, 93, 98] and attention-based [2, 11, 62, 65, 83] approaches have been proposed to encode temporal history. Some methods also explored the applicability of temporal convolutions [54, 66] and MLP-only architectures [10, 39] to the task of human motion forecasting. Additionally, various approaches have been proposed to model future human motion stochastically to produce diverse future sequence predictions, either with adversarial GAN formulations [7, 52], conditional variational autoencoders (VAEs) [3, 9, 12, 59, 63, 74, 87, 93, 95], or diverse sampling [21, 96]. More recently, diffusion methods [78, 79] have been used for human motion generation and forecasting [6, 19, 46, 84, 85, 94, 97, 99]. All of these methods require 3D ground truth sequences for training, which limits applicability to common scenarios where 3D inputs and ground-truth are not available. In contrast, ours requires only 2D training data for the action sequences, which is far more plentiful and easier to obtain. We are able to generate

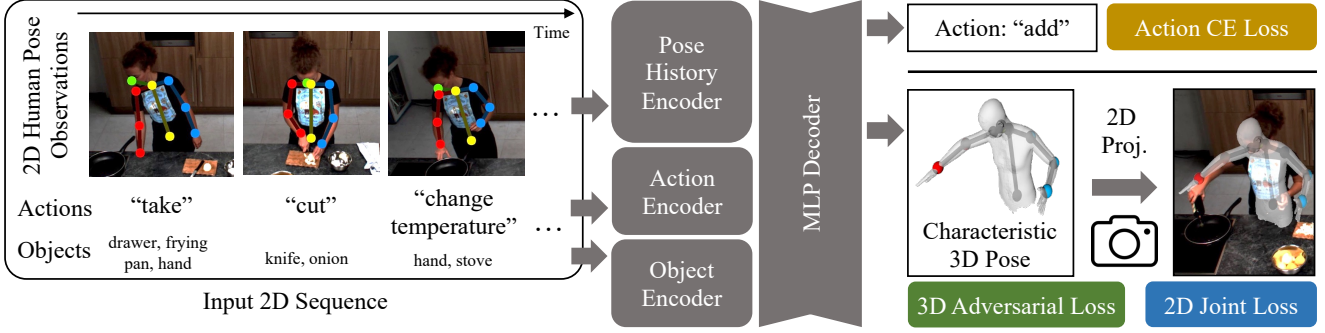


Figure 2. Our approach takes as input a sequence of RGB images, from which 2D poses are extracted, as well as their corresponding action label and initial set of objects. Each input is encoded into a joint latent space to jointly predict the next action label and characteristic 3D pose. While action labels are directly supervised, the 3D pose decoder is trained to match the next 2D action pose using differentiable projection, and an adversarial 3D loss encourages valid 3D pose prediction.

valid 3D poses by leveraging an adversarial loss formulation, operating on a database of uncorrelated 3D poses.

Human Action Forecasting. Action forecasting has been studied by many approaches to predict future actions from a sequence of observed actions [25, 26, 29, 48] or directly from an input video sequence [28, 32, 34, 35, 67, 75, 76, 76]. Various methods have been developed to learn effective representations, including Hidden Markov Models [51], RNNs [25–28, 32, 43, 76, 92], transformer-based networks [35, 36, 73], and self-supervised feature learning [41, 86]. There are approaches that focus on the short-term future [29, 32, 34, 35, 67, 73, 75, 76] or on longer-term actions [25, 26, 28, 29, 32, 34–36, 48, 67, 75, 76, 76]. Such method focus on characterizing anticipation with action labels only, while we aim to predict a richer characterization of the anticipated future by leveraging characteristic 3D poses, representative of future action goals in a sequence of action-pose predictions. Forecasting actions alongside human poses in 2D only has been studied in a few works, for 2D hand placement [57] or full-body 2D human poses at most 1 second into the future [100]. Our approach addresses the benefits of 3D reasoning in human motion forecasting, without requiring full 3D sequences for supervision.

Goal-Driven Future Prediction. Goal-driven forecasting has previously been explored beyond action label forecasting, and has been leveraged to predict goal locations for future human walking trajectories [14] and for future video sequences by predicting keyframes [5, 45, 70, 80]. Diller et al. [22] introduced the task of forecasting *characteristic 3D poses*, salient keyframe poses representing the next action. These goal-based poses are more semantically meaningful and consistent across different action sequences than time-based ones, and thus are better suited for long-term forecasting. We build upon these ideas by introducing a new goal-driven method for joint action anticipation and characteristic 3D pose forecasting in an auto-regressive system that can predict complex, long-term behavior sequences.

3. Method Overview

Our method aims to learn to jointly model future human actions along with the characteristic 3D poses representative of those actions. From a sequence of RGB image observations of a person performing a series of actions and the corresponding action labels, we predict a sequence of future action labels and 3D poses characteristic of these actions. This enables joint reasoning of not only global semantic behavior but also the physical manifestation thereof.

In the absence of 3D pose data of complex human actions, we weakly supervise forecasted 3D poses to align to future poses in 2D, and constrain the poses to be valid in 3D using an adversarial loss with a database of 3D poses. This does not require any correspondence between 3D pose data and 2D video, enabling action sequence supervision on commonly available 2D human action data together with carefully captured but unrelated human poses in 3D.

An overview of this approach is shown in Fig. 2. For an input sequence $S = \{(I_i, a_i, o_i)\}$ of N RGB images I_i with corresponding actions a_i and initially involved objects o_i , we aim to predict the future M actions $\{\hat{a}_k\}$ that will be taken along with their characteristic poses in 3D $\{\hat{Y}_k\}$. We define the human pose as a collection of J body joints at salient locations, so each output pose \hat{Y}_k is predicted as a set of J 3D coordinates. We first extract information about the observed 2D pose movement by detecting 2D poses $\{X_i\}$, each with J 2D joints, with a state-of-the-art 2D pose estimator that seamlessly integrates into our pipeline in a pre-trained and frozen form.

Next, we encode this information along with previously observed action and object labels to predict the next future action label \hat{a}_k and characteristic 3D pose \hat{Y}_k . We can then forecast a future sequence by autoregressively predicting a series, considering the 2D projections of the previously predicted 3D poses along with previously predicted actions as input to a new prediction.

4. Joint Forecasting of Actions and Characteristic 3D Poses

Our network takes as input the previous 2D observations $\{X_i\}$ extracted from the $\{I_i\}$ images, as well as action and object labels $\{a_i\}$ and $\{o_i\}$. Since we only predict action labels, object labels are given from the objects seen at the beginning of the sequence, and subsequently re-used for the entire sequence. Each of these are encoded in parallel with three separate encoders; the actions and objects with MLPs while the poses are projected into latent space with a single linear layer and then processed with a stack of three residual blocks. These encoded features are then all concatenated together in latent space, and processed jointly with an MLP to produce a common latent code z . Finally, we decode both poses and actions in parallel based on z using an MLP decoder each, yielding the next action label class as a vector $\hat{a}_k \in \mathbb{R}^{N_a}$ and 3D characteristic pose $\hat{Y}_k \in \mathbb{R}^{J \times 3}$, with N_a the number of action classes. For a more detailed architecture specification, we refer to the appendix.

We jointly learn future action labels and characteristic 3D poses by supervising \hat{a}_k and \hat{Y}_k to match the observed future 2D video, and constrain \hat{Y}_k to form a valid 3D pose by an adversarial loss, optimizing for the overall loss:

$$\mathcal{L} = \lambda_{action}\mathcal{L}_{action} + \lambda_{pose2d}\mathcal{L}_{pose2d} + \lambda_{adv3d}\mathcal{L}_{adv3d}, \quad (1)$$

where \mathcal{L}_{action} denotes the action loss, as described in Sec. 4.1, \mathcal{L}_{pose2d} and \mathcal{L}_{adv3d} constraining the predicted pose, as described in Sec. 4.2, and the λ weighting each loss.

4.1. Action Forecasting

Predicted future actions are decoded from the latent code z by an MLP decoder to predict the action class \hat{a}_k , supervised by cross entropy with the ground truth future action: $\mathcal{L}_{action} = \text{CE}(\hat{a}_k, a_k^{\text{gt}})$.

4.2. Characteristic Pose Forecasting

Our goal is to forecast complex action behavior not only in terms of action labels, but also manifested as a sequence of characteristic poses in 3D. Since we only have 2D pose annotations available, we first constrain these poses to represent future actions in 2D and make use of an adversarial regularization in 3D. This does not require any correspondence between 2D and 3D data, only a collection of valid 3D poses, which are readily available.

Differentiable 2D Projection

Our generator network predicts the next characteristic action pose \hat{Y}_k as a set of 3D joints. To constrain \hat{Y}_k based on the target future 2D pose X^{gt} extracted from the ground truth future image, we differentiably project \hat{Y}_k into the 2D

image with given camera parameter intrinsic K and extrinsic rotation and translation R, t :

$$\hat{X} = K(R\hat{Y}_k + t) \quad (2)$$

Since we learn from third-person video with a fixed camera, we can use the same camera parameters for all sequences used for training. We can then define the 2D pose loss as the mean squared error between the projected pose prediction and the ground truth:

$$\mathcal{L}_{pose2d} = \|X^{\text{gt}} - \hat{X}_k\|_2^2 \quad (3)$$

Note that we only predict the J joints that have been observed in the video data (excluding any joints that remain occluded in the observed video data), so this loss can be applied to all predicted joints.

Adversarial 3D Pose Regularization

While the action and pose prediction losses provide effective predictions when considered in the 2D projections, the \hat{Y}_k remain underconstrained in 3D and thus tend to exhibit large distortions and implausible bone lengths and angles, when trained with only 2D supervision. We thus constrain the predicted poses to form valid 3D poses by formulating an adversarial 3D loss from a critic network which that is simultaneously trained to distinguish predicted poses from a database of real 3D skeleton samples. Note that there is no correspondence between these skeletons and the 2D poses extracted from the action video sequences – any database of 3D skeletons can be used. We can thus train our approach with an entirely uncorrelated 3D pose dataset without requiring 3D pose annotation to action video.

We then formulate \mathcal{L}_{adv3d} as a Wasserstein loss [4], training the critic network in an alternating fashion with the generator. This enables effective forecasting of future 3D characteristic poses for predicted future action labels, without requiring any 3D observations as input.

In order to enable the critic network to learn effectively about likely intrinsic pose constraints (e.g., lengths, kinematic chains, or valid joint angles), the critic takes as input not only the 3D joint locations of \hat{Y}_k but also their kinematic statistics as a matrix Ψ , following [88, 89].

Ψ encodes joint angles and bone lengths as $\Psi = B^T B$, where $B = (b_1, b_2, \dots, b_b)$ is a matrix with columns $b_i = j_k - j_l$ representing the vectors between each joint j_k and j_l . Ψ then contains bone lengths l_i^2 on its diagonal, and angular representations on the off-diagonal entries.

4.3. Sequence Prediction

In order to forecast longer-term future behavior, our 3D pose predictions enable a natural autoregressive sequence prediction by taking the predictions \hat{X}_t, \hat{a}_t at time step t as

	MPII Cooking II				IKEA ASM			
	2d	3d	Action Accuracy		2d	3d	Action Accuracy	
Approach	MPJPE [px] ↓	Quality ↑	top-1 ↑	top-3 ↑	MPJPE [px] ↓	Quality ↑	top-1 ↑	top-3 ↑
Zero Velocity	118	–	–	–	74	–	–	–
Train Average	166	–	–	–	91	–	–	–
AVT [35] RGB	–	–	19%	42%	–	–	22%	49%
AVT [35] RGB+Skeleton	–	–	20%	40%	–	–	23%	47%
FUTR [36] RGB	–	–	27%	48%	–	–	19%	45%
FUTR [36] RGB+Skeleton	–	–	27%	49%	–	–	20%	46%
RepNet [88] + DLow [96]	89	0.70	–	–	51	0.30	–	–
RepNet [88] + GSPS [63]	75	0.57	–	–	55	0.11	–	–
RepNet [88] + STARS [93]	70	0.62	–	–	54	0.27	–	–
Joint 2D Pose & Action [100]	55	–	27%	43%	44	–	22%	46%
Ours	50	0.55	29%	51%	40	0.31	29%	50%

Table 1. Quantitative comparison with state-of-the-art action label and 3D pose forecasting. Our joint approach enables more accurate future action and pose predictions, compared to approaching both tasks separately, and outperforms joint action & 2D pose forecasting.

part of the input for time step $t + 1$. We can thus predict a sequence of M future action labels and characteristic 3D poses; we use $M = 10$ for MPII Cooking II [72] and $M = 5$ for IKEA-ASM [8], respectively.

4.4. Training Details

We train our approach for the $J = 9$ joints commonly seen across the input observed video data, characterizing the upper body in MPII Cooking II [72] and IKEA-ASM [8].

Additionally, we use loss weights $\lambda_{action} = 1e^6$, $\lambda_{pose} = 1$, and $\lambda_{adv3d} = 1$, empirically chosen to numerically balance each individual loss with the others.

We train our approach on a single NVIDIA GeForce RTX 2080TI for ≈ 12 hours until convergence. We use an ADAM optimizer with batch size 4096, weight decay 0.001, and a constant learning rate of 0.0001 for both generator and discriminator.

4.5. Datasets

We train and evaluate our approach on two datasets: MPII Cooking II [72] and IKEA-ASM [8]. Both datasets contain sequences of human actors performing complex, unscripted actions, and provide annotations of fine-grained sub-action steps. MPII Cooking II [72] is an action recognition dataset with 272 complex cooking sequences, each with an average of 35 annotated sub-actions. IKEA-ASM contains 370 sequences of actors assembling IKEA furniture, with an average of 15 annotated sub-actions.

In both datasets, each action sequence has been filmed from a fixed camera setup; the third-person point of view enables extraction of 2D poses with an off-the-shelf 2D pose estimator. We use OpenPose [13] in our experiments and note that our approach is agnostic to the concrete method of 2D pose detection. We provide more in-depth discussion and additional experiments in the appendix.

We consider the 9 upper-body joints of the OpenPose skeletons, as the other joints are almost always occluded in the video observations, and remove global translation by

centering each 2D pose at the neck joint.

Characteristic poses, in contrast to an arbitrary pose within a labeled action range, are the most representative pose of that action, and are annotated for all sub-actions in each sequence as the most articulated pose of that sub-action, following the annotation protocol of [22]. All annotations were performed by the authors within 5 days, yielding a total of $\approx 18,000$ characteristic poses ($\approx 12,000$ for MPII Cooking II and $\approx 6,000$ for IKEA-ASM). These poses are indicative of the action they represent as demonstrated in Tab. 2: Using such poses significantly improves performance, validating our annotation protocol.

For the 3D adversarial loss, we use $\approx 800,000$ human poses from popular 3D pose datasets: Human3.6m [42], AMASS [60], and GRAB [82]. Note that none of these 3D poses have any correspondence with the 2D posed actions from the MPII Cooking II dataset, instead depicting various human skeletons in natural and diverse poses.

5. Results

We evaluate sequence forecasting of action labels and characteristic 3D poses on the MPII Cooking II [72] and IKEA-ASM [8] datasets, and 3D pose quality by comparing to our database of high-fidelity 3D poses.

5.1. Evaluation Metrics

2D Pose Error. Since we only have 2D ground-truth data available for complex action sequences, we first project predicted 3D poses back into 2D, and evaluate the 2D mean per-joint position error (MPJPE) [42], in comparison with 2D poses extracted from ground-truth future frames using [13]: $E_{MPJPE} = \frac{1}{M} \sum_{j=1}^M \|\hat{X} - X^{gt}\|_2$.

3D Pose Quality. In the absence of annotated ground truth 3D poses for the action video sequences, we measure the quality of predicted 3D poses as how distinguishable they are in comparison to a set of real 3D poses. We follow [3] and evaluate quality by training a binary classifier on 50,000 human poses generated at different training steps




Figure 3. Action accuracy over time. Our joint action-characteristic pose forecasting enables more robust autoregressive action forecasting than action prediction without considering pose.

(representing examples of unrealistic 3D poses) and 50,000 real 3D pose samples. For classification accuracy a of this classifier, quality is measured as $1 - a$, with a quality of 1 indicating full indistinguishability from real poses. We refer to the appendix for more details on this quality metric.

Action Accuracy. We report the action accuracy of the predicted sequences, as the mean over all sequences in the test set. We evaluate the top- n accuracy based on whether the ground truth action is among the n highest scoring predictions, for $n = 1$ and $n = 3$.

5.2. Comparison to Human Pose Forecasting

Tab. 1 compares our method to state-of-the-art 3D pose forecasting methods DLow [96], GSPS [63], and STARS [93]. These methods expect sequences of observed 3D human poses as input; we thus first apply a state-of-the-art weakly supervised 3D pose estimator [88] on our 2D input poses, producing inputs and supervision in 3D. This method estimates 3D poses using an adversarial formulation, requiring a database of 3D poses not correlated with the 2D pose inputs. To ensure a fair comparison, this 3D pose database is exactly the same as the one our method uses.

We chose the 3D pose estimator of [88] since its weakly supervised formulation is most comparable to our approach. An additional comparison to a fully supervised approach for 3D pose lifting (SPIN [50]) is provided in the appendix.

We then train the 3D pose prediction methods from scratch on this generated data, using their original parameter settings. Each method takes as input a pose history of M poses and outputs a sequence of M poses, analogous to our setup where each pose is a characteristic pose corresponding to an action step ($M = 10$ for MPII Cooking II and $M = 5$ for IKEA-ASM). Our approach to lift 2D to future 3D poses and actions in an end-to-end fashion enables more effective pose forecasting than these state-of-the-art 3D pose forecasting approaches on both datasets.

In addition, we compare to the joint 2D action and pose forecasting approach of Zhu et al. [100]. Our approach

of forecasting long-term sequences of 3D poses alongside actions is able to outperform their 2D MPJPE pose prediction and action accuracy performance, due to improved spatial reasoning when forecasting 3D poses.

Statistical 2D Baselines. We additionally compare with two statistical baselines in 2D, following [22]: the average target train pose, and a zero-velocity baseline which was introduced by Martinez et al. [64] as competitive with state of the art. We outperform both baselines, indicating that our method learns a strong action pose representation.

5.3. Comparison to Action Label Forecasting

We compare the action accuracy of our joint action-pose forecasting to AVT [35] and FUTR [36], two state-of-the-art action anticipation methods, in Tab. 1. We train and evaluate both AVT and FUTR on input RGB frames and their action and object labels, equal to our training setup, and use their original training settings initialized with a pre-trained vision transformer [23] for AVT and extracted I3D features [15] from our datasets for FUTR. Additionally, as we consider extracted 2D poses from the input RGB images, we also evaluate a variant that is trained and evaluated on RGB images overlaid with 2D poses (+Skeleton). Our approach outperforms these baselines in both scenarios, by jointly predicting future actions and characteristic 3D poses.

5.4. Ablation Studies

What is the effect of pose forecasting on long-term action understanding? Tab. 3 shows that there is a notable improvement in action accuracy between training only with an action loss vs. training action and 2D pose loss jointly. This becomes more apparent when training action only vs action and full pose prediction (2D and 3D losses). In addition, Fig. 3 shows the correspondence between autoregressive prediction length and action accuracy: jointly forecasting poses and actions is enables more robust autoregressive forecasting over time. We conclude that pose forecasting is beneficial for long-term action understanding.

How does action forecasting affect pose prediction performance? Tab. 3 demonstrates that pose forecasting trained jointly with action prediction is complementary and enables more accurate pose prediction.

Poses		2D	3D	Action Accuracy	
Train	Test	MPJPE [px] ↓	Quality ↑	top-1 ↑	top-3 ↑
Uncoupled	Uncoupled	75	0.29	28%	48%
Middle	Middle	58	0.45	26%	43%
Random	Random	67	0.37	22%	42%
Characteristic	Characteristic	50	0.55	29%	51%

Table 2. Ablation on pose forecasting on MPII Cooking II [72]. Our characteristic pose representation maximizes MPJPE and action performance: We consider pose prediction following state-of-the-art pose forecasting as decoupled from actions (uncoupled), as well as poses coupled to actions but in the middle of an action range, or at a random time therein, and our characteristic pose prediction. The same pose type is used for both train and evaluation.

			MPII Cooking II				IKEA ASM			
Losses During Training			2D	3D	Action Accuracy		2D	3D	Action Accuracy	
Action	2D Proj.	3D Adv.	MPJPE [px] ↓	Quality ↑	top-1 ↑	top-3 ↑	MPJPE [px] ↓	Quality ↑	top-1 ↑	top-3 ↑
✓	×	×	—	—	21%	41%	—	—	24%	45%
✓	✓	×	62	0.10	26%	49%	46	0.05	27%	49%
×	✓	×	54	0.21	—	—	44	0.09	—	—
×	✓	✓	58	0.53	—	—	43	0.29	—	—
✓	✓	✓	50	0.55	29%	51%	40	0.31	29%	50%

Table 3. Ablation on the effect of the action, 2D projection, and 3D adversarial losses. Combining all together for joint forecasting enables complementary learning to produce the best performance.

What is the effect of characteristic pose forecasting?

Since state-of-the-art pose forecasting focuses on fixed frame rate predictions independent of actions, we compare with such joint forecasting of action and pose where predicted poses are sampled at equally spaced points in time in Tab. 2 (uncoupled). Additionally, we consider alternative poses to forecast for each action rather than a characteristic 3D pose (middle of the annotated action range, and randomly selected within the action range). We keep the same pose representation for training and testing (i.e., evaluate on middle poses when trained on them, etc.), for a fair comparison. We observe the best performance when forecasting characteristic 3D poses along with action labels, showing their usefulness for forecasting long sequences of 3D poses and actions.

5.5. Qualitative Results

Qualitative evaluations for the predicted poses are shown in Fig. 5 on data from MPII Cooking II [72] and in Fig. 4 on data from IKEA-ASM [8]. We compare our approach with

state-of-the-art 3D pose forecasting of DLow [96], GSPS [63], and STARS [93]. For each method, we show a 3D body mesh in addition to the predicted 3D pose joints, to more comprehensively show the 3D structure of the forecasting results; we obtain body meshes by fitting SMPL [58] to each methods’ predicted 3D body joints.

As there is no 3D ground truth available, we show the camera perspective with background for context as well as without background for a 3D pose only version. The two views demonstrate the fit to the ground truth 2D along with the quality of the 3D pose, respectively. Our approach leads to poses that better follow the ground-truth action poses in 2D compared to both previous methods while still maintaining a valid pose structure in 3D. Notably, this is true for both datasets, as our approach effectively forecasts the different data characteristics of both cooking as well as furniture assembly. In particular, our joint action-3D pose forecasting enables more accurate forecasting with diverse and accurate 3D pose structures.

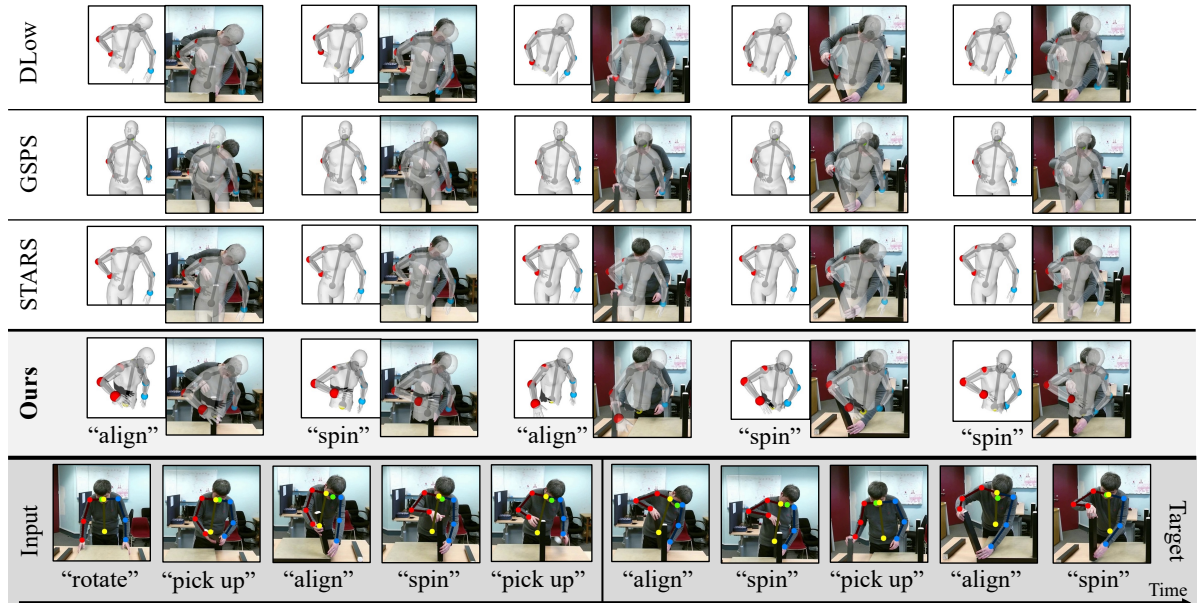


Figure 4. Qualitative comparison between DLow [96], GSPS [63], STARS [93], and our method on IKEA-ASM [8] data. For each method, we show the 3D predicted pose projected into the 2D target view, without background (small) and with background for context (full size). Our joint reasoning captures the individual characteristic action poses more faithfully while producing spatially plausible 3D poses.

5.6. Limitations

While we have demonstrated the potential of joint action and 3D pose forecasting, several limitations remain. For instance, our method leverages a separate 2D pose extraction as input to training, while an end-to-end formulation could potentially better leverage other useful signal in the input frames. Additionally, a more holistic body representation than pose joints would be important for finer-grained interactions that involve reasoning over small limbs (e.g., hands) and body surface contact.

6. Conclusion

In this paper, we proposed to forecast future human behavior by jointly predicting future actions alongside characteristic 3D poses. We do not require any 3D annotated action sequences, or 3D input data; instead, we learn complex action sequences from 2D action video data, and regularize predicted poses with an adversarial formulation against uncorrelated 3D pose data. Experiments demonstrate that our joint forecasting enables complementary feature learning, outperforming each individual task considered separately.

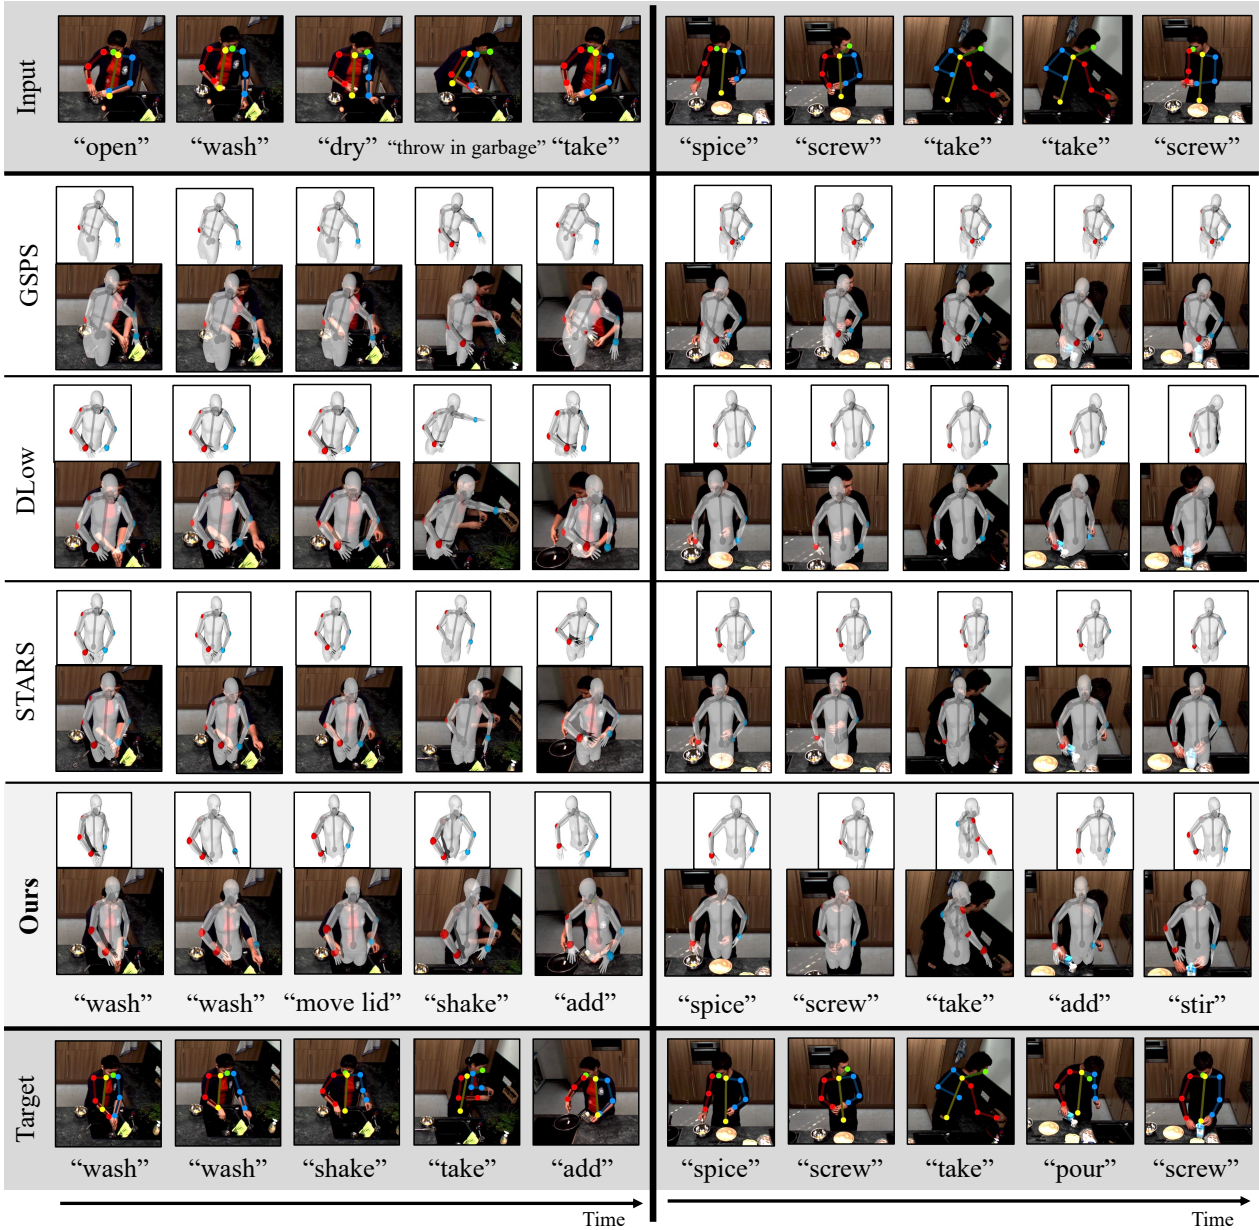


Figure 5. Qualitative comparison between DLow [96], GSPS [63], STARS [93], and our method on two sequences (left and right) from MPII Cooking II [72]. For each method, we show the 3D predicted pose projected into 2D, without background (small) and with background for context (full size). By considering both 3D pose and action forecasting together, we more effectively forecast the longer-term behavior.

Acknowledgements

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), and the German Research Foundation (DFG) Grant “Learning How to Interact with Scenes through Part-Based Understanding”.

References

- [1] Emre Aksan, Manuel Kaufmann, and Otmar Hilliges. Structured prediction helps 3d human motion modelling. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 7143–7152. IEEE, 2019. 2
- [2] Emre Aksan, Manuel Kaufmann, Peng Cao, and Otmar Hilliges. A spatio-temporal transformer for 3d human motion prediction. In *International Conference on 3D Vision, 3DV 2021, London, United Kingdom, December 1-3, 2021*, pages 565–574. IEEE, 2021. 2
- [3] Sadeqh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Lars Petersson, and Stephen Gould. A stochastic conditioning scheme for diverse human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5223–5232, 2020. 2, 5
- [4] Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 214–223. PMLR, 2017. 4
- [5] Amir Bar, Roei Herzig, Xiaolong Wang, Anna Rohrbach, Gal Chechik, Trevor Darrell, and Amir Globerson. Compositional video synthesis with action graphs. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, pages 662–673. PMLR, 2021. 3
- [6] Germán Barquero, Sergio Escalera, and Cristina Palmero. Belfusion: Latent diffusion for behavior-driven human motion prediction. *CoRR*, abs/2211.14304, 2022. 2
- [7] Emad Barsoum, John Kender, and Zicheng Liu. Hp-gan: Probabilistic 3d human motion prediction via gan. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 1418–1427, 2018. 2
- [8] Yizhak Ben-Shabat, Xin Yu, Fatemeh Saleh, Dylan Campbell, Cristian Rodriguez-Opazo, Hongdong Li, and Stephen Gould. The ikea asm dataset: Understanding people assembling furniture through actions, objects and pose. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 847–859, 2021. 2, 5, 7, 14, 15, 16, 17
- [9] Andreas Blattmann, Timo Milbich, Michael Dorkenwald, and Björn Ommer. Behavior-driven synthesis of human dynamics. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 12236–12246. Computer Vision Foundation / IEEE, 2021. 2
- [10] Arij Bouazizi, Adrian Holzbock, Ulrich Kressel, Klaus Dietmayer, and Vasileios Belagiannis. Motionmixer: Mlp-based 3d human body pose forecasting. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 791–798. ijcai.org, 2022. 2
- [11] Yujun Cai, Lin Huang, Yiwei Wang, Tat-Jen Cham, Jianfei Cai, Junsong Yuan, Jun Liu, Xu Yang, Yiheng Zhu, Xiaohui Shen, Ding Liu, Jing Liu, and Nadia Magnenat-Thalmann. Learning progressive joint propagation for human motion prediction. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part VII*, pages 226–242. Springer, 2020. 2
- [12] Yujun Cai, Yiwei Wang, Yiheng Zhu, Tat-Jen Cham, Jianfei Cai, Junsong Yuan, Jun Liu, Chuanxia Zheng, Sijie Yan, Henghui Ding, Xiaohui Shen, Ding Liu, and Nadia Magnenat-Thalmann. A unified 3d human motion synthesis model via conditional variational auto-encoder*. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 11625–11635. IEEE, 2021. 2
- [13] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 5, 15, 17
- [14] Zhe Cao, Hang Gao, Karttikeya Mangalam, Qi-Zhi Cai, Minh Vo, and Jitendra Malik. Long-term human motion prediction with scene context. In *European Conference on Computer Vision*, pages 387–404. Springer, 2020. 3
- [15] João Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 4724–4733. IEEE Computer Society, 2017. 6, 15
- [16] Hsu-Kuang Chiu, Ehsan Adeli, Borui Wang, De-An Huang, and Juan Carlos Niebles. Action-agnostic human pose forecasting. In *IEEE Winter Conference on Applications of Computer Vision, WACV 2019, Waikoloa Village, HI, USA, January 7-11, 2019*, pages 1423–1432. IEEE, 2019. 2
- [17] Qiongjie Cui and Huaijiang Sun. Towards accurate 3d human motion prediction from incomplete observations. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 4801–4810. Computer Vision Foundation / IEEE, 2021. 2
- [18] Qiongjie Cui, Huaijiang Sun, and Fei Yang. Learning dynamic relationships for 3d human motion prediction. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 6518–6526. Computer Vision Foundation / IEEE, 2020. 2
- [19] Rishabh Dabral, Muhammad Hamza Mughal, Vladislav Golyanik, and Christian Theobalt. Mofusion: A framework for denoising-diffusion-based motion synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 9760–9770. IEEE, 2023. 2

- [20] Lingwei Dang, Yongwei Nie, Chengjiang Long, Qing Zhang, and Guiqing Li. MSR-GCN: multi-scale residual graph convolution networks for human motion prediction. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 11447–11456. IEEE, 2021. 2
- [21] Lingwei Dang, Yongwei Nie, Chengjiang Long, Qing Zhang, and Guiqing Li. Diverse human motion prediction via gumbel-softmax sampling from an auxiliary space. In *MM '22: The 30th ACM International Conference on Multimedia, Lisboa, Portugal, October 10 - 14, 2022*, pages 5162–5171. ACM, 2022. 2
- [22] Christian Diller, Thomas Funkhouser, and Angela Dai. Forecasting characteristic 3d poses of human actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15914–15923, 2022. 2, 3, 5, 6
- [23] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. 6
- [24] Haoshu Fang, Jiefeng Li, Hongyang Tang, Chao Xu, Haoyi Zhu, Yuliang Xiu, Yong-Lu Li, and Cewu Lu. Alphapose: Whole-body regional multi-person pose estimation and tracking in real-time. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(6):7157–7173, 2023. 14
- [25] Yazan Abu Farha and Juergen Gall. Uncertainty-aware anticipation of activities. In *2019 IEEE/CVF International Conference on Computer Vision Workshops, ICCV Workshops 2019, Seoul, Korea (South), October 27-28, 2019*, pages 1197–1204. IEEE, 2019. 3
- [26] Yazan Abu Farha, Alexander Richard, and Juergen Gall. When will you do what? - anticipating temporal occurrences of activities. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 5343–5352. Computer Vision Foundation / IEEE Computer Society, 2018. 3
- [27] Yazan Abu Farha, Alexander Richard, and Juergen Gall. When will you do what? - anticipating temporal occurrences of activities. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 5343–5352. Computer Vision Foundation / IEEE Computer Society, 2018.
- [28] Yazan Abu Farha, Qiuhong Ke, Bernt Schiele, and Juergen Gall. Long-term anticipation of activities with cycle consistency. In *Pattern Recognition - 42nd DAGM German Conference, DAGM GCPR 2020, Tübingen, Germany, September 28 - October 1, 2020, Proceedings*, pages 159–173. Springer, 2020. 3
- [29] Basura Fernando and Samitha Herath. Anticipating human actions by correlating past with the future with jaccard similarity measures. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 13224–13233. Computer Vision Foundation / IEEE, 2021. 3
- [30] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 4346–4354. IEEE Computer Society, 2015. 2
- [31] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In *Proceedings of the IEEE international conference on computer vision*, pages 4346–4354, 2015. 2
- [32] Antonino Furnari and Giovanni Maria Farinella. What would you expect? anticipating egocentric actions with rolling-unrolling lstms and modality attention. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 6251–6260. IEEE, 2019. 3
- [33] Antonino Furnari and Giovanni Maria Farinella. Rolling-unrolling lstms for action anticipation from first-person video. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):4021–4036, 2020. 2
- [34] Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. Predicting the future: A jointly learnt model for action anticipation. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 5561–5570. IEEE, 2019. 3
- [35] Rohit Girdhar and Kristen Grauman. Anticipative video transformer. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 13485–13495. IEEE, 2021. 2, 3, 5, 6, 14
- [36] Dayoung Gong, Joonseok Lee, Manjin Kim, Seong Jong Ha, and Minsu Cho. Future transformer for long-term action anticipation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 3042–3051. IEEE, 2022. 2, 3, 5, 6, 14
- [37] Anand Gopalakrishnan, Ankur Mali, Dan Kifer, C. Lee Giles, and Alexander G. Ororbia II. A neural temporal model for human motion prediction. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 12116–12125. Computer Vision Foundation / IEEE, 2019. 2
- [38] Liang-Yan Gui, Yu-Xiong Wang, Xiaodan Liang, and José M. F. Moura. Adversarial geometry-aware human motion prediction. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part IV*, pages 823–842. Springer, 2018. 2
- [39] Wen Guo, Yuming Du, Xi Shen, Vincent Lepetit, Xavier Alameda-Pineda, and Francesc Moreno-Noguer. Back to MLP: A simple baseline for human motion prediction. In *IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2023, Waikoloa, HI, USA, January 2-7, 2023*, pages 4798–4808. IEEE, 2023. 2
- [40] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social GAN: socially acceptable trajectories with generative adversarial networks. In *2018 IEEE*

- Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 2255–2264. Computer Vision Foundation / IEEE Computer Society, 2018. 14
- [41] Tengda Han, Weidi Xie, and Andrew Zisserman. Memory-augmented dense predictive coding for video representation learning. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part III*, pages 312–329. Springer, 2020. 3
- [42] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(7):1325–1339, 2014. 2, 5, 17
- [43] Ashesh Jain, Avi Singh, Hema S Koppula, Shane Soh, and Ashutosh Saxena. Recurrent neural networks for driver activity anticipation via sensory-fusion architecture. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3118–3125. IEEE, 2016. 3
- [44] Ashesh Jain, Amir Roshan Zamir, Silvio Savarese, and Ashutosh Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 5308–5317. IEEE Computer Society, 2016. 2
- [45] Dinesh Jayaraman, Frederik Ebert, Alexei A. Efros, and Sergey Levine. Time-agnostic prediction: Predicting predictable video frames. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. 3
- [46] Chiyu "Max" Jiang, Andre Cornman, Cheolho Park, Benjamin Sapp, Yin Zhou, and Dragomir Anguelov. Motion-diffuser: Controllable multi-agent motion prediction using diffusion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 9644–9653. IEEE, 2023. 2
- [47] Jigyasa Singh Katroliya, Ahmed El-Sherif, Hartmut Feld, Bruno Mirbach, Jason R. Rambach, and Didier Stricker. Ticam: A time-of-flight in-car cabin monitoring dataset. In *32nd British Machine Vision Conference 2021, BMVC 2021, Online, November 22-25, 2021*, page 277. BMVA Press, 2021. 14
- [48] Qiuhong Ke, Mario Fritz, and Bernt Schiele. Time-conditioned action anticipation in one shot. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 9925–9934. Computer Vision Foundation / IEEE, 2019. 3
- [49] Kris M Kitani, Brian D Ziebart, James Andrew Bagnell, and Martial Hebert. Activity forecasting. In *European conference on computer vision*, pages 201–214. Springer, 2012. 2
- [50] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *ICCV*, 2019. 6, 15
- [51] Hilde Kuehne, Ali Arslan, and Thomas Serre. The language of actions: Recovering the syntax and semantics of goal-directed human activities. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 780–787, 2014. 2, 3
- [52] Jogendra Nath Kundu, Maharshi Gor, and R. Venkatesh Babu. Bihmp-gan: Bidirectional 3d human motion prediction GAN. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 8553–8560. AAAI Press, 2019. 2
- [53] Bin Li, Jian Tian, Zhongfei Zhang, Hailin Feng, and Xi Li. Multitask non-autoregressive model for human motion prediction. *IEEE Trans. Image Process.*, 30:2562–2574, 2021. 2
- [54] Chen Li, Zhen Zhang, Wee Sun Lee, and Gim Hee Lee. Convolutional sequence to sequence model for human dynamics. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 5226–5234. Computer Vision Foundation / IEEE Computer Society, 2018. 2
- [55] Maosen Li, Siheng Chen, Yangheng Zhao, Ya Zhang, Yanfeng Wang, and Qi Tian. Dynamic multiscale graph neural networks for 3d skeleton based human motion prediction. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 211–220. Computer Vision Foundation / IEEE, 2020. 2
- [56] Maosen Li, Siheng Chen, Zihui Liu, Zijing Zhang, Lingxi Xie, Qi Tian, and Ya Zhang. Skeleton graph scattering networks for 3d skeleton-based human motion prediction. In *IEEE/CVF International Conference on Computer Vision Workshops, ICCVW 2021, Montreal, BC, Canada, October 11-17, 2021*, pages 854–864. IEEE, 2021. 2
- [57] 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 *European Conference on Computer Vision*, pages 704–721. Springer, 2020. 3
- [58] 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, 2015. 7
- [59] Thomas Lucas, Fabien Baradel, Philippe Weinzaepfel, and Grégory Rogez. Posegpt: Quantization-based 3d human motion generation and forecasting. In *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part VI*, pages 417–435. Springer, 2022. 2
- [60] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *International Conference on Computer Vision*, pages 5442–5451, 2019. 2, 5, 17
- [61] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion

- prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9489–9497, 2019. [2](#)
- [62] Wei Mao, Miaomiao Liu, and Mathieu Salzmann. History repeats itself: Human motion prediction via motion attention. In *European Conference on Computer Vision*, pages 474–489. Springer, 2020. [2](#)
- [63] Wei Mao, Miaomiao Liu, and Mathieu Salzmann. Generating smooth pose sequences for diverse human motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13309–13318, 2021. [2](#), [5](#), [6](#), [7](#), [8](#), [14](#), [15](#), [16](#), [17](#)
- [64] Julieta Martinez, Michael J. Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 4674–4683. IEEE Computer Society, 2017. [2](#), [6](#)
- [65] Ángel Martínez-González, Michael Villamizar, and Jean-Marc Odobez. Pose transformers (POTR): human motion prediction with non-autoregressive transformers. In *IEEE/CVF International Conference on Computer Vision Workshops, ICCVW 2021, Montreal, BC, Canada, October 11-17, 2021*, pages 2276–2284. IEEE, 2021. [2](#)
- [66] Omar Medjaouri and Kevin Desai. HR-STAN: high-resolution spatio-temporal attention network for 3d human motion prediction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2022, New Orleans, LA, USA, June 19-20, 2022*, pages 2539–2548. IEEE, 2022. [2](#)
- [67] Antoine Miech, Ivan Laptev, Josef Sivic, Heng Wang, Lorenzo Torresani, and Du Tran. Leveraging the present to anticipate the future in videos. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 2915–2922. Computer Vision Foundation / IEEE, 2019. [3](#)
- [68] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019. [17](#)
- [69] Dario Pavllo, David Grangier, and Michael Auli. Quaternet: A quaternion-based recurrent model for human motion. In *British Machine Vision Conference 2018, BMVC 2018, Newcastle, UK, September 3-6, 2018*, page 299. BMVA Press, 2018. [2](#)
- [70] Karl Pertsch, Oleh Rybkin, Jingyun Yang, Shenghao Zhou, Konstantinos Derpanis, Kostas Daniilidis, Joseph Lim, and Andrew Jaegle. Keyframing the future: Keyframe discovery for visual prediction and planning. In *Learning for Dynamics and Control*, pages 969–979. PMLR, 2020. [3](#)
- [71] Nicholas Rhinehart and Kris M Kitani. First-person activity forecasting with online inverse reinforcement learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3696–3705, 2017. [2](#)
- [72] Marcus Rohrbach, Anna Rohrbach, Michaela Regneri, Sikandar Amin, Mykhaylo Andriluka, Manfred Pinkal, and Bernt Schiele. Recognizing fine-grained and composite activities using hand-centric features and script data. *International Journal of Computer Vision*, pages 1–28, 2015. [2](#), [5](#), [6](#), [7](#), [8](#), [14](#), [16](#), [17](#)
- [73] Debaditya Roy and Basura Fernando. Action anticipation using pairwise human-object interactions and transformers. *IEEE Trans. Image Process.*, 30:8116–8129, 2021. [3](#)
- [74] Tim Salzmann, Marco Pavone, and Markus Ryhl. Motron: Multimodal probabilistic human motion forecasting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 6447–6456. IEEE, 2022. [2](#)
- [75] Fadime Sener and Angela Yao. Zero-shot anticipation for instructional activities. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 862–871. IEEE, 2019. [3](#)
- [76] Fadime Sener, Dipika Singhania, and Angela Yao. Temporal aggregate representations for long-range video understanding. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XVI*, pages 154–171. Springer, 2020. [3](#)
- [77] Theodoros Sofianos, Alessio Sampieri, Luca Franco, and Fabio Galasso. Space-time-separable graph convolutional network for pose forecasting. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 11189–11198. IEEE, 2021. [2](#)
- [78] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 2256–2265. JMLR.org, 2015. [2](#)
- [79] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [2](#)
- [80] Didac Suris, Ruoshi Liu, and Carl Vondrick. Learning the predictability of the future. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 12607–12617. Computer Vision Foundation / IEEE, 2021. [3](#)
- [81] Wandu Susanto, Marcus Rohrbach, and Bernt Schiele. 3d object detection with multiple kinects. In *European Conference on Computer Vision*, pages 93–102. Springer, 2012. [17](#)
- [82] Omid Taheri, Nima Ghorbani, Michael J. Black, and Dimitrios Tzionas. Grab: A dataset of whole-body human grasping of objects. In *Computer Vision – ECCV 2020*, pages 581–600, Cham, 2020. Springer International Publishing. [2](#), [5](#), [17](#)
- [83] Yongyi Tang, Lin Ma, Wei Liu, and Wei-Shi Zheng. Long-term human motion prediction by modeling motion context and enhancing motion dynamics. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 935–941. ijcai.org, 2018. [2](#)

- [84] Julian Tanke, Linguang Zhang, Amy Zhao, Chengcheng Tang, Yujun Cai, Lezi Wang, Po-Chen Wu, Juergen Gall, and Cem Keskin. Social diffusion: Long-term multiple human motion anticipation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9601–9611, 2023. [2](#)
- [85] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. [2](#)
- [86] Carl Vondrick, Hamed Pirsiavash, and Antonio Torralba. Anticipating visual representations from unlabeled video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 98–106, 2016. [3](#)
- [87] Jacob Walker, Kenneth Marino, Abhinav Gupta, and Martial Hebert. The pose knows: Video forecasting by generating pose futures. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 3352–3361. IEEE Computer Society, 2017. [2](#)
- [88] Bastian Wandt and Bodo Rosenhahn. Repnet: Weakly supervised training of an adversarial reprojection network for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7782–7791, 2019. [4](#), [5](#), [6](#), [14](#), [15](#), [16](#), [17](#)
- [89] Bastian Wandt, Hanno Ackermann, and Bodo Rosenhahn. 3d reconstruction of human motion from monocular image sequences. *IEEE transactions on pattern analysis and machine intelligence*, 38(8):1505–1516, 2016. [4](#)
- [90] Borui Wang, Ehsan Adeli, Hsu-Kuang Chiu, De-An Huang, and Juan Carlos Niebles. Imitation learning for human pose prediction. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 7123–7132. IEEE, 2019. [2](#)
- [91] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. [14](#)
- [92] Yu Wu, Linchao Zhu, Xiaohan Wang, Yi Yang, and Fei Wu. Learning to anticipate egocentric actions by imagination. *IEEE Transactions on Image Processing*, 30:1143–1152, 2020. [3](#)
- [93] Sirui Xu, Yu-Xiong Wang, and Liang-Yan Gui. Diverse human motion prediction guided by multi-level spatial-temporal anchors. In *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXII*, pages 251–269. Springer, 2022. [2](#), [5](#), [6](#), [7](#), [8](#), [14](#), [15](#), [16](#), [17](#)
- [94] Sirui Xu, Zhengyuan Li, Yu-Xiong Wang, and Liang-Yan Gui. Interdiff: Generating 3d human-object interactions with physics-informed diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14928–14940, 2023. [2](#)
- [95] Xinchun Yan, Akash Rastogi, Ruben Villegas, Kalyan Sunkavalli, Eli Shechtman, Sunil Hadap, Ersin Yumer, and Honglak Lee. Mt-vae: Learning motion transformations to generate multimodal human dynamics. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 265–281, 2018. [2](#)
- [96] Ye Yuan and Kris Kitani. Dlow: Diversifying latent flows for diverse human motion prediction. In *European Conference on Computer Vision*, pages 346–364. Springer, 2020. [2](#), [5](#), [6](#), [7](#), [8](#), [14](#), [15](#), [16](#), [17](#)
- [97] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *CoRR*, abs/2208.15001, 2022. [2](#)
- [98] Chongyang Zhong, Lei Hu, Zihao Zhang, Yongjing Ye, and Shihong Xia. Spatio-temporal gating-adjacency GCN for human motion prediction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 6437–6446. IEEE, 2022. [2](#)
- [99] Zixiang Zhou and Baoyuan Wang. UDE: A unified driving engine for human motion generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 5632–5641. IEEE, 2023. [2](#)
- [100] Yanjun Zhu, David Doermann, Yanxia Zhang, Qiong Liu, and Andreas Girgensohn. What and how? jointly forecasting human action and pose. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 771–778. IEEE, 2021. [3](#), [5](#), [6](#)

Appendix

We show in this appendix additional qualitative and quantitative results (Sec. A and Sec. B), demonstrate the ability of our method to generalize to multi-actor scenarios (Sec. C), verify our method’s robustness to 2D detection results (Sec. D), detail our baseline evaluation protocol (Sec. E), elaborate on the 3D quality metric (Sec. F), show the architecture used in our approach (Sec. G), and provide additional details regarding the data (Sec. H).

A. Additional Qualitative Results

Fig. 6 shows additional qualitative results of our method, on both MPII Cooking 2 [72] (left column) and IKEA-ASM [8] (right column), as compared to pose baselines DLow [96], GSPS [63], and STARS [93].

B. Additional Quantitative Results

B.1. Characteristic Poses

Analogous to Tab. 2 in the main paper, Tab. 6 shows an ablation on pose timings and compares our approach of using characteristic poses to poses taken at regular time intervals (“uncoupled”) as well as in the middle or at a random time of an action, on IKEA-ASM [8] data. To further illustrate this point, Tab. 4 shows an additional ablation with poses at random points in the sequence, but at most 1s from the closest characteristic pose. This already improves performance compared to other approaches while still being outperformed by directly using characteristic poses.

B.2. Input Noise Ablation

Tab. 5 shows the effect using a noise vector as additional input to our method. It encourages more diversity in predictions, which benefits pose and action forecasting.

Poses	2D	3D	Action Accuracy	
	MPJPE [px] ↓	Quality ↑	top-1 ↑	top-3 ↑
Uncoupled	75	0.29	28%	48%
Middle	58	0.45	26%	43%
Random	67	0.37	22%	42%
Centered on Char. Poses	69	0.33	28%	50%
Characteristic	50	0.55	29%	51%

Table 4. Ablation on pose forecasting on MPII Cooking II [72]. We consider pose prediction following state-of-the-art pose forecasting as decoupled from actions (uncoupled), as well as poses coupled to actions in various fashions: middle (the middle pose of an action range), random (a random pose of the action), random but at most 1s from the closest characteristic pose (centered), and our characteristic pose prediction.

B.3. Input Objects Ablation

Object inputs slightly improve results (Tab. 5), due to added context for broad actions like “add,” e.g. “add ingredient” vs. “add water to pot.”.

C. Multi-Actor Interaction Scenario

In addition to our experiments with single human actors in the main paper, we show here that our approach is able to generalize to multi-actor scenarios, with minor modifications. We show this in Tab. 7 with additional dataset TICaM [47] where driver and passenger are interacting in an in-car driving scenario (actions include “talking”, various handoffs). Our modifications are: (1) Additional encoder and decoder for the second person (2) Interaction pooling introduced in Social GAN [40]. Our modified method outperforms simple combinations of previous works, with and without interaction modelling, demonstrating the wide applicability of our method.

D. 2D Input Pose Quality

In Fig. 8, we replace OpenPose with AlphaPose [24] and Detectron2 [91], both only slightly changing the final results, indicating that our method does not depend on a specific 2D pose detector. We also experiment with added random noise to OpenPose: our method remains relatively robust. The coupled changes in pose and action accuracy further demonstrate the effectiveness of our joint feature learning.

E. Baseline Evaluation Details

E.1. State-of-the-Art Pose Forecasting

We evaluate the performance of our baselines using the same input data that is available to our method. Pose forecasting baselines DLow [96], GSPS [63], and STARS [93] are trained and evaluated on sequences of our manually annotated characteristic poses. Since there is no ground-truth 3D pose data available, we first use RepNet [88], a state-of-the-art 3D pose estimation method, to retrieve 3D skeletons from our 2D characteristic poses. We train this method from scratch using *the same* database of valid 3D poses that is available to our method, allowing for a fair comparison.

E.2. State-of-the-Art Action Label Forecasting

We train action baselines AVT [35] and FUTR [36] using sequences of our characteristic pose frames together with the corresponding action labels as input. For AVT, we use their default parameters used by the original authors for their ablation on third-person dataset 50Salads/Breakfast, inputting our RGB frames instead. For a fair comparison, we also supply the action and object history for each step by encoding both label sequences with a small encoder (a single linear layer) each and fuse these features with the


	MPII Cooking II				IKEA ASM			
	2d	3d	Action Accuracy		2d	3d	Action Accuracy	
Approach	MPIPE [px] ↓	Quality ↑	top-1 ↑	top-3 ↑	MPIPE [px] ↓	Quality ↑	top-1 ↑	top-3 ↑
No Objects	61	0.52	28%	51%	42	0.30	29%	50%
No Noise	55	0.49	29%	50%	48	0.29	30%	51%
Ours	50	0.55	29%	51%	40	0.31	29%	50%

Table 5. Ablations studies with no object input and no noise input.

	2D	3D	Action Accuracy	
Poses	MPIPE [px] ↓	Quality ↑	top-1 ↑	top-3 ↑
Uncoupled	64	0.30	28%	48%
Middle	47	0.35	28%	47%
Random	49	0.24	28%	49%
Characteristic	41	0.35	29%	50%

Table 6. Ablation on pose forecasting, on the IKEA-ASM [8] dataset. We consider predicting poses following state-of-the-art pose forecasting in a decoupled fashion from actions (uncoupled), as well as poses coupled to actions in various fashions: middle (the middle pose of an action range), random (a random pose of the action), and our characteristic pose prediction, which benefits action prediction the most.

	2d	3d	Action Accuracy	
Approach	MPIPE [px] ↓	Quality ↑	top-1 ↑	top-3 ↑
FUTR RGB + Skeleton	-	-	38%	64%
RepNet + STARS	89	0.34	-	-
Ours (No Interactions)	68	0.40	40%	67%
Ours (Interaction Modeling)	58	0.41	48%	73%



Setting

Table 7. Multi-Actor Scenario on TICA dataset (Singh et al. 21).

MPII Cooking II	2d	3d	Action Accuracy	
Approach	MPIPE [px] ↓	Quality ↑	top-1 ↑	top-3 ↑
OpenPose + max. 20px noise	59	0.45	26%	47%
OpenPose + max. 10px noise	57	0.47	26%	46%
Ours (using Detectron2)	47	0.54	28%	55%
Ours (using AlphaPose)	46	0.57	28%	56%
Ours (using OpenPose)	50	0.55	29%	51%

Table 8. Robustness to 2D Pose Input Quality & Detectors.

image features generated by the AVT encoder. For FUTR, we first generate I3D features [15] from our RGB frames and concatenate them with action and object history after encoding these in the same way as for AVT.

We then train two variants of both methods: One with the raw RGB frames, action history, and object history as

input (“AVT RGB” and “FUTR RGB” in the main results figure), and one with additional 2D skeleton input (skeletons rendered on top of the RGB frames) from the skeletons that we extract with OpenPose [13] (“AVT RGB+Skeleton” and “FUTR RGB+Skeleton”).

E.3. Supervised 3D Pose Lifting

For better comparability, we used weakly supervised approach [88] for pose lifting. This is important, since there is no ground-truth coupling between 2D and corresponding 3D action poses in our setting. Nevertheless, we compare to baselines [63, 93, 96] in Tab. 9 with poses lifted using fully supervised pre-trained SPIN [50]; our approach outperforms even these improved baselines in terms of 2D MPIPE.

	MPII Cooking II		IKEA ASM	
	2d	3d	2d	3d
Approach	MPIPE [px] ↓	Quality ↑	MPIPE [px] ↓	Quality ↑
SPIN [50] + DLow [96]	81	0.89	43	0.43
SPIN [50] + GSPS [63]	74	0.66	45	0.29
SPIN [50] + STARS [93]	66	0.80	41	0.40
Ours	50	0.55	40	0.31

Table 9. Comparison to pose baselines using fully-supervised pre-trained 3D pose estimation method SPIN [50]. In our main experiments, we instead compare to weakly supervised baseline RepNet [88] for a fair comparison.

F. 3D Quality Metric Details

For our pose quality metric, we use a 3-layer MLP binary classifier of 3D poses. Training poses are randomly sampled from ground-truth (real) and predicted (fake) poses from each training epoch, producing a total of 100k real and fake poses each. Fake poses exhibit a range of small to large unrealistic deformations, depending on when they were sampled, ranging from random joint placements to widely inconsistent bone lengths to unnatural joint angles to only minor inconsistencies in the bone lengths.

As an additional intuitive metric we show the mean absolute bone length difference of right and left body in 3D in Tab. 10. We observe that this metric correlates with our classifier-based quality.

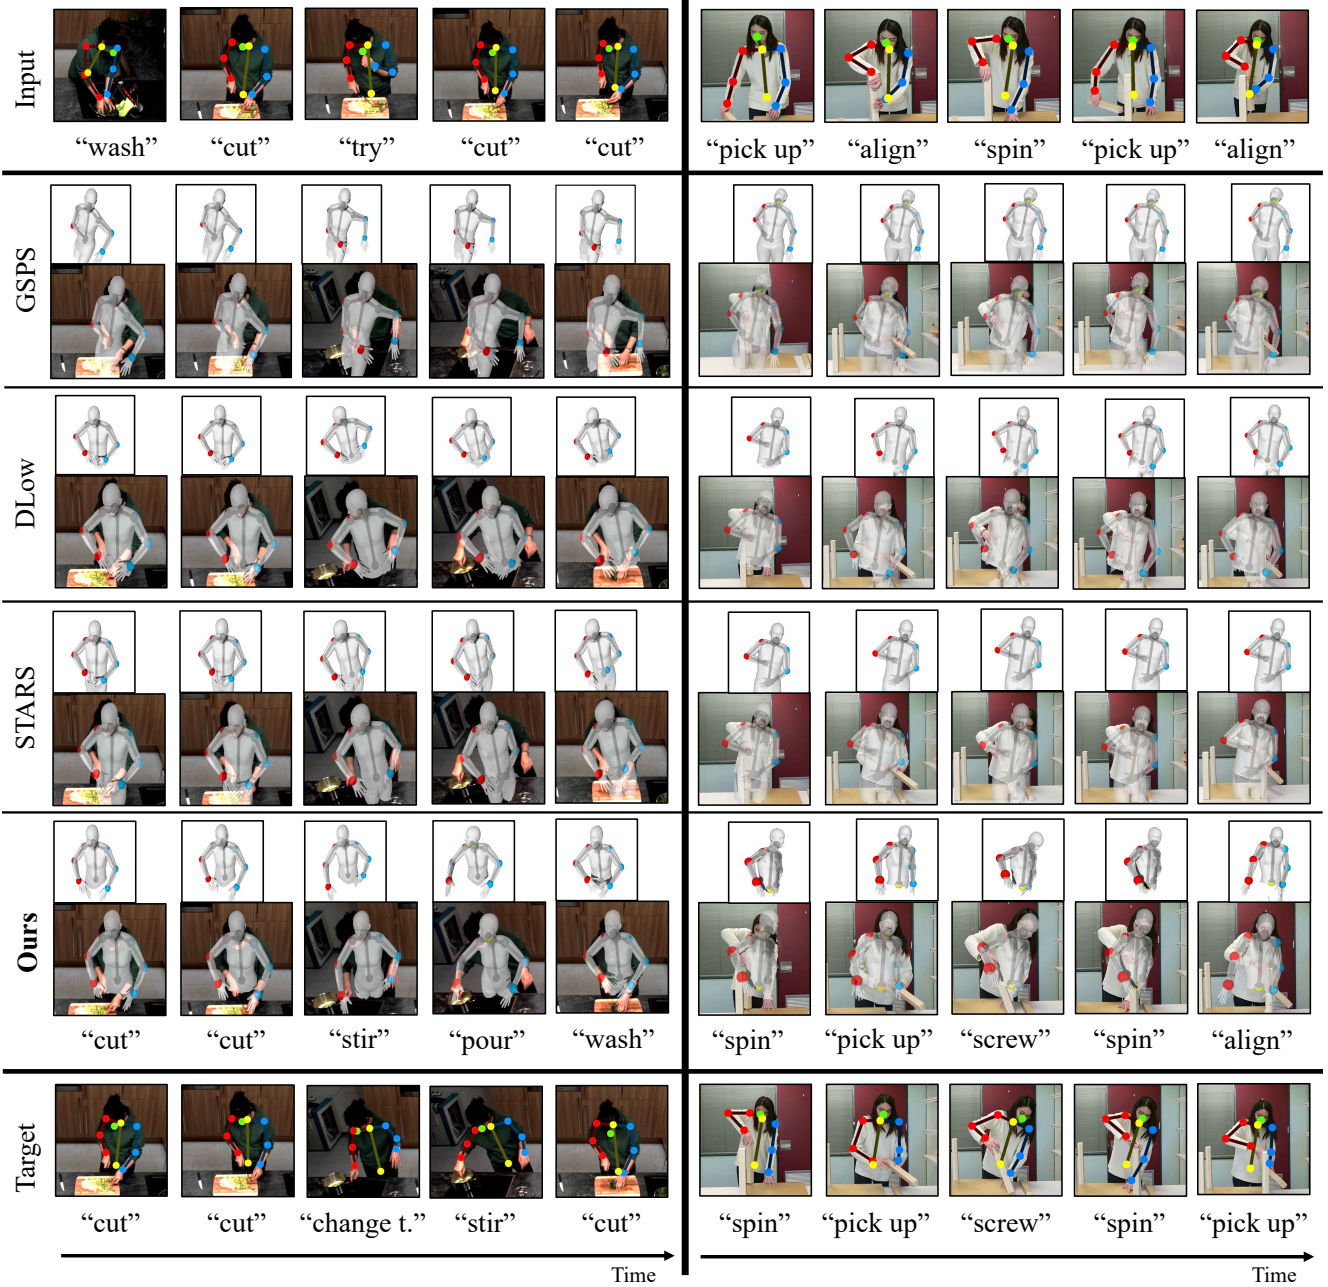


Figure 6. Additional qualitative comparison between DLoW [96], GSPS [63], STARS [93], and our method on two sequences (left on MPII Cooking 2 [72], right on IKEA-ASM [8]). For each method, we show a the 3D predicted pose projected into the 2D target view, without background for a pose only version (small) as well as with background for context (full size).

G. Architecture Details

Generator Network Fig. 7 shows our generator architecture in detail with input and output dimensions for linear layers, and the slope for leaky ReLU layers.

Critic Network Our adversarial critic network processes generator outputs with 4 linear layers and 3 kinematic chain layers which are designed to encourage correct bone lengths

(as shown in [88]), in parallel. 2 linear layers then combine both outputs and produce the final critic score.

H. Data Details

H.1. Camera Parameters

While intrinsic camera parameters are often available in captured image data, the camera parameters for captured

Approach	MPII Cooking II		IKEA ASM	
	Symm. [mm] ↓	Quality ↑	Symm. [mm] ↓	Quality ↑
RepNet [88] + DLow [96]	14	0.70	45	0.30
RepNet [88] + GSPS [63]	21	0.57	59	0.11
RepNet [88] + STARS [93]	16	0.62	46	0.27
No 3D Adversarial Loss	75	0.10	66	0.05
2D Projection Loss Only	57	0.21	61	0.09
No Action Loss	22	0.53	39	0.29
Ours	22	0.55	39	0.31

Table 10. Additional quality metric and its correlation to our classifier-based metric: Absolute bone length difference between right and left body, compared to pose baselines and ablations.

video were not available from the MPII Cooking 2 [72] dataset to use for pose projection. We thus optimized for intrinsic camera parameters from the video sequence data in correspondence with the 3D scene reconstruction of the empty kitchen environment, as given by [81].

For IKEA-ASM [8], we use the provided intrinsic camera parameters directly.

H.2. Pose Joint Layout

We use the 9 upper-body joints of the native OpenPose [13] joint layout for skeletons in 2D, and adapt skeletons in our 3D database to use the same format. Tab. 11 shows the correspondence between our joint layout, OpenPose [13], Human3.6m [42], and SMPL-X [68]. 3D datasets AMASS [60] and GRAB [82] provide human bodies in SMPL-X format; we first extract their skeleton joints using their publicly available code and then convert it into our layout using the correspondences in Tab. 11.

H.3. MPII Cooking 2 Details

We use action labels as annotated in the 2D cooking action dataset MPII Cooking 2 [72]. These annotations provide action labels (87 classes) for frame ranges in each sequence as well as the involved objects (187 classes). We first cluster similar actions together, yielding a total of 37 action clusters, which we then use as action classes in our experiments.

In addition, since our goal is to forecast upper-body actions with objects in the foreground, we remove instances of poses and corresponding actions that occur in the background - e.g., when taking out objects from the cupboard, or from the fridge.

In total, there are 272 cooking action sequences; we create a random train/val/test split along sequences with a ratio of 70% / 15% / 15%, yielding 190, 40, 40 sequences for each set.

H.4. IKEA-ASM Details

We use action labels as annotated in the IKEA furniture assembly dataset IKEA-ASM [8]. These annotations provide action labels (31 classes) for frame ranges in each sequence; we use them without explicit object information since each action already encodes its associated object.

In total, there are 370 furniture assembly action sequences; we create a random train/val/test split along sequences with a ratio of 70% / 15% / 15%, yielding 227, 48, 48 sequences for each set.

Ours		OpenPose		Human3.6m		SMPL-X	
Idx	Name	Idx	Name	Idx	Name	Idx	Name
0	head	0	nose	15	head	15	head
1	neck	1	neck	13	thorax	12	neck
2	right shoulder	2	right shoulder	25	right shoulder	17	right shoulder
3	right elbow	3	right elbow	26	right elbow	19	right elbow
4	right hand	4	right hand	27	right wrist	42	right index 3
5	left shoulder	5	left shoulder	17	left shoulder	16	left shoulder
6	left elbow	6	left elbow	18	left elbow	18	left elbow
7	left hand	7	left wrist	19	left wrist	27	left index 3
8	hip	8	mid-hip	0	hip	0	pelvis

Table 11. Joint layout used in our experiments, for both 2D and 3D skeletons.

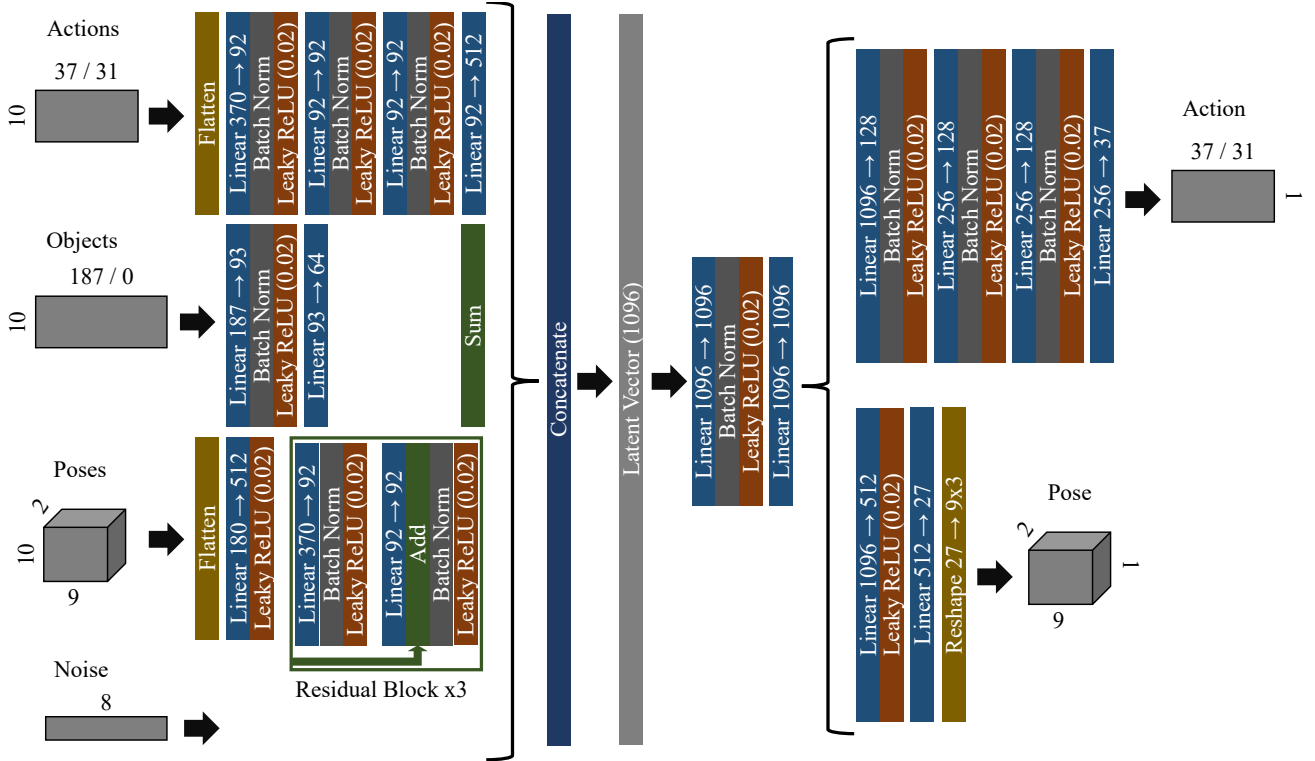


Figure 7. Network architecture specification.