IMoS: Intent-Driven Full-Body Motion Synthesis for Human-Object Interactions

Anindita Ghosh^{1,2,3} Rishabh Dabral^{2,3} Vladislav Golyanik^{2,3} Christian Theobalt^{2,3} Philipp Slusallek^{1,3}

¹ German Research Center for Artificial Intelligence (DFKI) ²Max-Planck Institute or Informatics ³Saarland Informatics Campus



Figure 1: Visualizations of Motion Sequences Showing Virtual Characters Performing Various Intended Actions with Different Objects, as Generated by Our Method. We synthesize the full-body pose sequences along with the 3D object positions from textual inputs. Our method can synthesize single-handed as well as two-handed interactions depending on the intent and the type of the object used.

Abstract

Can we make virtual characters in a scene interact with their surrounding objects through simple instructions? Is it possible to synthesize such motion plausibly with a diverse set of objects and instructions? Inspired by these questions, we present the first framework to synthesize the full-body motion of virtual human characters performing specified actions with 3D objects placed within their reach. Our system takes as input textual instructions specifying the objects and the associated 'intentions' of the virtual characters and outputs diverse sequences of full-body motions. This is in contrast to existing works, where full-body action synthesis methods generally do not consider object interactions, and human-object interaction methods focus mainly on synthesizing hand or finger movements for grasping objects. We accomplish our objective by designing an intent-driven full-body motion generator, which uses a pair of decoupled conditional variational autoencoders (CVAE) to learn the motion of the body parts in an autoregressive manner. We also optimize for the positions of the objects with six degrees of freedom (6DoF) such that they plausibly fit within the hands of the synthesized characters. We compare our proposed method with the existing methods of motion synthesis and establish a new and stronger state-of-the-art for the task of intent-driven motion synthesis. Through a user study, we further show that our synthesized full-body motions appear more realistic to the participants in more than 80% of scenarios compared to the current state-of-the-art methods, and are perceived to be as good as the ground truth on several occasions. Project page: vcai.mpi-inf.mpg.de/IMoS

© 2023 The Author(s)

Computer Graphics Forum © 2023 The Eurographics Association and John Wiley & Sons Ltd. Published by John Wiley & Sons Ltd.

1. Introduction

Humans regularly use and interact with objects in numerous ways in the real world. Interactions like eating a fruit or brushing the teeth, as shown in Fig. 1, are part of our daily-routines. Being able to synthesize such interactions in a *virtual* 3D environment through textual instructions has widespread applications in several areas, including computer graphics and robotics [ALNM20], movie script visualization [HMLC09] and game design [SSR07]. For instance, in a digitally created movie scene or a virtual role-playing game, a character might need to interact with the scene objects based on a set of instructions, such as yielding tools, using objects, or eating various items. Manually modeling such 3D character-object interactions or *intentions* is time-consuming and laborious, in particular when we desire to synthesize a variety of possible motions with the same intention and object.

In this context, many recent methods automatically synthesize motions for virtual characters by encoding control signals such as music [LYL*19; LRX*21; LYC*20], speech [BCRM21] or text, either as a speech transcript [BRB*21] or as a highlevel action descriptions [AM19; LWC*18; AHC*18] or sentences [GCO*21; PBV22; GZW*20]. Methods synthesizing fullbody pose sequences typically follow an autoregressive approach to maintain continuity in the synthesized motions [LZCvdP20; RBH*21; GZZ*22]. These autoregressive motion synthesis frameworks predict short-term future sequences from a short history. We can also find several methods for hand-object interactions [JLWW21; CKA*22; KYZ*20; TGBT20; ZYSK21], which focus on generating only the wrist and finger movements for grasping various objects. However, to create a plausible motion sequence for an intent-driven virtual character, modeling hand-motion alone is not sufficient. Rather, we believe it is crucial to operate in the space of full-body motion synthesis. There are two prime reasons for this. Firstly, synthesizing full body movements allow for a broader range of interactions (Fig. 1). For several intents like eating, drinking, exchanging objects between hands, inspecting, and passing, the head, the arms, and the torso are also part of the complete action sequence [TGBT20]. Secondly, trivially attaching the synthesized hand motion to the remaining body [PRB*18] leads to an uncanny and physically implausible motion generation (see suppl. video). Further, recent works [TCBT22; WWZ*22] have demonstrated the ability to generate whole-body grasping motion starting from a T-Pose till the moment of the grasp. However, synthesizing a plausible motion sequence after the first grasp moment, especially based on an intent guiding the human-object interaction remains unaddressed. To address this gap in the literature, we propose a novel framework to synthesize diverse sequences of plausible body motions to animate different character-object interactions given input textual instructions consisting of actions (intentions) and objects (Fig. 2).

We learn generalizable intent encodings from the input intentobject pairs using a CLIP encoder [RKH*], which is a large-scale language model trained on a large corpus of text-image pairs. Given the initial body poses and the 3D object positions, we design an intent-driven full-body motion generator model to autoregressively generate full-body motions (Sec.3). We follow a decoupling approach and model the arms and the body motions using separate INPUT

OUTPUT



Figure 2: Overview of Our Intent-Driven Full-Body Motion Generator. Our model takes in initial 3D body poses and object positions (upper-left), and instruction labels (upper-middle) describing the object types and the intended actions. We design a pair of decoupled conditional variational autoencoders or CVAEs, the Arms Syntheis Module and the Body Synthesis Module (lowermiddle), to separately synthesize the arms and the rest of the body. We also design a Condition Encoder (middle) to condition our decoupled CVAEs based on the input instruction labels and the body shape parameters. We concatenate our synthesized arm and body motions, and use our Object Optimizer Module (lower-right) to optimize for the 6-DoF parameters of the object such that it satisfies grasping constraints. Our model outputs the synthesized full-body motion sequence along with object positions (upper-right).

Conditional Variational Auto-Encoders (CVAE) [SLY15] to make our output arm and body movements more precise. Since the autoencoders are variational in nature, they allow us to sample diverse motions from the latent space at inference time. We also observe that regressing the motion from a longer history is crucial in modeling long-term temporal dependence between the joints. We use a position-encoded self-attention mapping to model correlations between the different joints to allow a broader range of interactions. Finally, we perform an optimization routine to estimate the corresponding 6-DoF object positions, relative to the hand position in each frame (Sec 3.2.4). We also use the recovered object positions to condition the future motion synthesis.

We train and evaluate our method on the recently introduced GRAB dataset [TGBT20] (Sec. 5.1), consisting of $\sim 1.3K$ sequences of human-object interactions displaying multiple intents. We quantitatively evaluate our synthesized sequences on metrics like the Mean Per Joint Position Error, the Average Variance Error, and the Average Pairwise Distance to test the variability of the model. More importantly, we conduct a visual perceptual study for subjective evaluation of our synthesized motions compared to recent conditional motion synthesis methods (Sec. 5.6).

In summary, our primary technical contributions are threefold:

• A new framework for generating diverse motion sequences in which virtual humanoid characters interact with objects of

	Motion Synthesis					
Method	Full Body	Intent- Driven	Only Till Grasp	Object Manipulation		
GRABNet [TGBT20]	×	×	×	×		
D-Grasp [CKA*22]	×	×	×	✓		
A2M [PBV21]	1	 Image: A second s	×	×		
ACTOR [PBV21]	1	1	×	×		
GOAL [TCBT22]	1	×	1	×		
SAGA [WWZ*22]	1	×	1	×		
Ours	 Image: A second s	 Image: A second s	×	✓		

Table 1: **Overview of the Problem Definitions of Existing Methods.** Our approach is the only one combining three important characteristics, and the first one to synthesize intent-driven full-body pose sequences for motions with object manipulation.

known shapes, placed within their reach, from text-based instruction labels. In contrast to previous works on characterobject interactions, our proposed method also optimizes for the 6-DoF object positions in 3D.

- We place special focus on synthesizing interactions involving *both* hands. This also includes sequences where the object is exchanged between the hands ("offhand") a setting that has not been addressed before.
- To achieve this, we learn separate variational latent embeddings for the arms and the rest of the body. These latent embeddings are conditioned jointly on the types of objects used as well as the intended actions to be performed with them. This enables diversity in the synthesized motions and allows the accurate synthesis of both-handed interactions.

2. Related Work

We position our work with respect to the past works that model human-object interactions in 3D. We study them from the following four vantage points: human pose forecasting and synthesis, humanobject 3D interaction modeling, hand-object grasp synthesis, and full-body grasp synthesis.

Human Pose Forecasting and Synthesis. Human pose forecasting methods predict future motions from a sequence of past poses as joint positions [MBR17], joint rotations [PGA18], or even SMPL body models [RBH*21]. Recent works on 3D human pose forecasting are stochastic methods [LLW*21; YK20] that use Variational Autoencoders [KW14] or GANs [GPM*20] to bring some variability in the output motion sequences. HuMoR [RBH*21] proposes a CVAE architecture that learns a distribution of pose transitions in the latent space while also ensuring physical plausibility through a post-processing optimization. Motion-VAE [LZCvdP20] learns to drive a character based on a goal position by decoding from a variational latent space. Characteristic 3D pose [DFD22] uses a probabilistic approach to predict a future 3D characteristic pose given a short sequence of observations.

Human motion synthesis models, on the other hand, are trained to synthesize a motion sequence conditioned either on semantic action labels [PBV21; GZW*20; DFD22], or text sen-

© 2023 The Author(s) Computer Graphics Forum © 2023 The Eurographics Association and John Wiley & Sons Ltd. tences [PBV22; GZZ*22]. Action2Motion [GZW*20] inputs an action label to generate the human pose in an autoregressive manner using a VAE-GRU. Differently, ACTOR [PBV21] employs a VAE-Transformer [VSP*17] to generate the full sequence in one shot. TEMOS [PBV22] uses the VAE-Transformer concept on a multimodal setting to generate motions from text sentences. Our goal is to not only synthesize full-body poses depending on a semantic label, but also to take into account object interactions.

Human-Object 3D Interaction Modeling. With the availability of several human-object 3D datasets like [SZKS19], BE-HAVE [BXP*22], PROX [HCTB19], D3D-HOI [XJMS21], H2O [KTS*21], GraviCap [DSJ*21], joint human-object motion modeling is an active research topic. In this paradigm, PHOSA [ZPJ*20] reconstructs the human and the object in the scene by jointly optimizing for the reprojection error of the object and the human. Neural State Machines [SZKS19] synthesizes human motion while interacting with objects like chairs or a wall in the scene. Likewise, SAMP [HCV*21] incorporates a path planning module to improve the character's motion in the scene. We, on the other hand, work deal with synthesizing fine-grained motions with handheld objects using instruction labels as inputs.

Hand-Object Grasp Synthesis. Grasp synthesis has been extensively studied in computer graphics [KP15; ES03; KYZ*20; ZYSK21; LFP07] and robotics [DKB*10; HL06; LPX*19; AGK18; BI05]. Some have approached this analytically by formulating grasp synthesis to be a constrained optimization problem satisfying the properties of a grasp [KDCI10; SKK12]. Several data-driven approaches [PG16; RA15] focus on learning the representations for synthesizing grasps through machine learning methods. More relevant recent approaches [KYZ*20; JLWW21; TGBT20; BHHF19] predict the hand parameters of the MANO hand model [RTB17] for synthesizing a grasp using neural networks. A large number of image datasets [BTT*20; LWM21; ZLM*19; ZHT*21; HVT*19] featuring hand-object interaction with contact maps are currently available. Taheri et al. [TGBT20] further introduce the GRAB dataset, which captures not only the contact map from hand, but also the full-body motion before and during the grasp. They also propose GrabNet, a network that estimates MANO parameters at the moment of grasp for unseen objects in a coarse-to-fine manner. [KYZ*20] proposes Grasping Field, a method that learns an implicit representation of the handobject interaction using a generative model. Grady et al. [GTT*21] derive physically plausible hand pose estimation by optimizing estimated hand meshes with contact prediction. We differ from all the methods in that our work focuses on synthesizing full-body sequences. While modeling hand-object interaction is an interesting problem, it is inherently limited in its ability to model several human-object interaction types which manifest themselves only when the full human body is accounted for (e.g., tilting back the head when drinking from a glass).

Full-Body Grasp Synthesis. GOAL [TCBT22], proposed recently, synthesizes full-body motion for grasping a given object. They first estimate the whole-body grasping pose for a given object and treat this pose as their goal for a motion infilling module which interpolates the motion between a T-Pose and the goal pose.

A. Ghosh, R. Dabral, V. Golyanik, C. Theobalt, P. Slusallek / IMoS: Intent-Driven Full-Body Motion Synthesisfor Human-Object Interactions



Figure 3: Architecture of Our Intent-Driven Full-Body Motion Generator Model. Given previous k frames of body poses and object positions, we train the arms and the rest of the body separately using our Arm Synthesis (upper-middle) and the Body Synthesis (lower-middle) Modules respectively. We design both these models using conditional variational autoencoders. We jointly synthesize the entire motion sequences autoregressively, conditioned on the input intent, the object used, and the body shape parameter, all encoded through our Condition Encoder (upper-left). We use position-encoded self-attention on the past k frames for the body joints before passing them through our Body Synthesis Module. After generating the body pose, our Object Optimizer Module (lower-right) optimizes for the 6-DoF pose of the given object such that it plausibly fits within the hands of the synthesized character. Please refer to Sec. 3.2 for more details.

SAGA [WWZ*22] also follows a similar strategy of motion infilling but uses markers to represent body pose while also learning a contact map for the grasp for additional supervision. Both these methods synthesize full-body motions *until* the point of grasping. Our method differs from these existing methods (see Table 1), since we synthesize the motion taking place *after* the object is grasped. We believe this is non-trivial and a more challenging setup. Conditioning human and object motions on a certain interaction type, while also ensuring diversity in the generated motion sequences, requires additionally learning their intent-based mutual interactions in an efficient and generalizable manner.

3. Intent-Driven Full-Body Motion Generator

We show the architecture of our intent-driven full-body motion generator model in Fig. 3. Given a human character's shape and initial 3D body pose, a rigid 3D object placed within their reach, and an intended action to perform with that object, our goal is to synthesize a full-body motion sequence of the character grasping the object and performing the intended action. We pose this problem as synthesizing the full-body motion sequence conditioned on a textual instruction label indicating the *intent* of the subject and the said object. We solve this problem through four main modules. First, we encode the input instruction labels consisting of the type of the object and the associated action using our Condition Encoder. We also input the subject's body shape parameters into our Condition Encoder. We use this encoding as a conditioning signal for the rest of the modules.

A key characteristic of our problem setting is that the arms are the primary movers during human-object interactions. Therefore, we use a pair of decoupled CVAEs to separately synthesize the arm movements and the rest of the body movements, using an Arm Synthesis Module and a Body Synthesis Module, respectively. Finally, we use an Object Optimizer Module to optimize for the 6-DoF pose of the given object such that it plausibly fits within the hands of the synthesized character.

3.1. 3D Human Body and Object Representation

We represent the human mesh using the SMPL-X [PCG*19] parametric body model. SMPL-X parametrizes the full human body along with hands and faces as a differentiable function $SMPLX(\beta, \mathbf{r}, \Psi, \mathbf{t})$ of the body shape parameters $\beta \in \mathbb{R}^{10}$, the root translation $\mathbf{t} \in \mathbb{R}^3$, the axis-angle rotations for the body joints $\mathbf{r} \in \mathbb{R}^{J \times 3}$ (J = 55), and the face expression parameters $\Psi \in \mathbb{R}^{10}$. It maps the parameters to a body mesh with 10,475 vertices. To improve the stability and the convergence characteristics of our model, we use the 6D continuous representations [ZBL*19] $\theta \in \mathbb{R}^{J \times 6}$ to represent body joint rotations. We downsample all the objects in the dataset to 300 vertices for faster optimization. The object's 6-DOF pose is represented using a rotation matrix $\mathbf{R} \in \mathbb{R}^9$ and a translation vector $\mathbf{T} \in \mathbb{R}^3$.

3.2. Model Design

We now discuss each of our modules in detail. Our synthesis pipeline assumes that the human interacts with only one object at a time. Interactions can be either one-handed or both-handed depending on the type of action and the shape of the object.

3.2.1. Condition Encoder

We input the object's category label using a one-hot vector $\mathbf{w_o} \in \mathbb{R}^{51}$. To represent the intended action information, we pass the intended action label, given as an English word, through the pre-trained CLIP [TGH*22] model and use the embeddings $\mathbf{w_a} \in \mathbb{R}^{512}$ that it outputs. The idea behind encoding the action labels with a

pre-trained text encoder is the general relevance between the action semantics and the corresponding body movements. For example, actions such as "eating" and "drinking" typically invoke similar body movements and are also semantically close, whereas other actions such as "taking pictures" and "offhanding" invoke different body movement and are also semantically different. Therefore, their corresponding embeddings given by a large-scale language model such as CLIP provides a regularized, semantic-based distribution of the intended actions and stabilizes further processing.

We concatenate $\mathbf{w_o}$ and $\mathbf{w_a}$ with the body shape parameters ($\beta \in \mathbb{R}^{10}$) and pass them into our Condition Encoder $\mathbf{q_c}$. Our Condition Encoder uses a series of MLPs to encode these input signals and projects them onto an encoded feature vector $\phi \in \mathbb{R}^{400}$ as

$$\phi = \mathbf{q_c}(\mathbf{w_o}, \mathbf{w_a}, \beta). \tag{1}$$

3.2.2. Arm Synthesis Module

Our Arm Synthesis Module uses a conditional variational autoencoder that synthesizes the arm movements, conditioned on our condition encoder output ϕ and the previous *k* frames of synthesized arm poses along with the 3D object positions. The encoder of this module, $\mathbf{q}_{\mathbf{a}}$, takes in the tuple $\mathbf{q}_{\mathbf{a}}{}^{in} = \{\phi, \theta^a_{\mathbf{t}-\mathbf{k}:\mathbf{t}-1}, \mathbf{T}_{\mathbf{t}-\mathbf{k}:\mathbf{t}-1}, \mathbf{R}_{\mathbf{t}-\mathbf{k}:\mathbf{t}-1}\}$, where $\theta^a_{\mathbf{t}-\mathbf{k}:\mathbf{t}-1}$ are the rotations for the arm joints synthesized by the past *k* frames, and $\mathbf{T}_{\mathbf{t}-\mathbf{k}:\mathbf{t}-1}, \mathbf{R}_{\mathbf{t}-\mathbf{k}:\mathbf{t}-1}$ are the translation and rotation parameters of the object for the past *k* frames. During training, $\mathbf{q}_{\mathbf{a}}$ uses a series of MLPs on the inputs and maps them to the parameters of a latent normal distribution, $\mu_{\mathbf{a}}, \sigma_{\mathbf{a}} \in \mathbf{R}^{32}$. The decoder, $\hat{\mathbf{q}}_{\mathbf{a}}$, samples $\mathbf{z}_{\mathbf{a}} \in \mathbb{R}^{32}$ from the latent distribution and uses the previous pose information ($\hat{\mathbf{q}}_{\mathbf{a}}^{\mathbf{in}}$) to synthesize the arm pose for the current frame ($\hat{\theta}_{\mathbf{t}}^{\mathbf{a}}$) through a series of MLPs with skip connections as

$$\hat{\theta}_{\mathbf{t}}^{\mathbf{a}} = \hat{\mathbf{q}}_{\mathbf{a}}(\mathbf{z}_{\mathbf{a}}, \mathbf{q}_{\mathbf{a}}^{\mathbf{m}}). \tag{2}$$

3.2.3. Body Synthesis Module

Similar to our Arm Synthesis Module, our Body Synthesis module is based on a conditional variational autoencoder. We use the term 'body' to denote the rest of the body parts apart from the arms, including the head, the torso, the hips, and the legs. We also note that the movements of all these parts are correlated when performing a full-body action. For example, to drink from a cup, one has to tilt their head back when bringing the cup to their mouth.

To model such fine-grained correlations, we first compute a selfattention mapping between all the joints in each pose as

$$\boldsymbol{\theta}_{\mathbf{k}}^{\mathbf{pe}} = \left[Attn(\mathbf{Q}, \mathbf{K}, \mathbf{V})\right]_{\mathbf{k}},\tag{3}$$

where the query **Q** is a joint position and the key-value pair (**K**, **V**) are information of all other joints provided as *J* sinusoidal positional encodings for each of the *k* frames. The encoder of the module, **q**_b, takes in the tuple **q**_b^{*in*} = $\left\{\phi, \hat{\theta}^a_t, \Theta^{\text{pe}}_{t-k:t-1}, \mathbf{T}_{t-k:t-1}, \mathbf{R}_{t-k:t-1}\right\}$. The structure of **q**_b is similar to that of the Arm Synthesis Module encoder **q**_a, and it maps the input **q**ⁱⁿ_b to the parameters of a latent normal distribution, $\mu_b, \sigma_b \in \mathbf{R}^{100}$. The decoder, $\hat{\mathbf{q}}_b$, samples $\mathbf{z}_b \in \mathbb{R}^{100}$ from the latent





Figure 4: **Our Hand-Object Setup.** We design the energy term E_d to enforce that the distances between the hand and the object vertices remain constant throughout the synthesis. Through the hand-object contact term E_c , we also enforce that the points in contact in the first frame remain in contact during the synthesis.

distribution and outputs the rest of the body poses as

$$\hat{\theta}_{\mathbf{t}}^{\mathbf{b}} = \hat{\mathbf{q}}_{\mathbf{b}}(\mathbf{z}_{\mathbf{b}}, \mathbf{q}_{\mathbf{b}}^{\mathbf{in}}). \tag{4}$$

We then concatenate $\hat{\theta}_t^a$ and $\hat{\theta}_t^b$ to obtain the full-body pose $\hat{\theta}_t$ at time *t*. We pass $\hat{\theta}_t$ to our Object Optimizer Module, along with the last predicted object position, to generate the object position for the current frame.

3.2.4. Object Optimizer Module

We have so far focused only on synthesizing the body poses for a given instruction. For a complete synthesis, we also need to estimate the corresponding 6-DoF positions of the object. Although fine-grained object synthesis is not the main goal of our work, we aim to produce plausible object trajectories faithful to the synthesized full-body motion. To this end, our core assumptions are that (*a*) at the moment of grasping in the initial frame, the object is at rest in an upright position and (*b*) inter-vertex distances between the vertices of the object and the hand remain constant throughout our intent-driven motion synthesis.

With these assumptions, we optimize for the object's rotation \mathbf{R} , translation \mathbf{T} , as well as the pose parameters of the hand, \mathbf{P}^h , in the SMPL-X parameter space.

We first compute the matrix of Euclidean distances $\mathbf{D} \in \mathbb{R}^{N \times M}$ between the vertices in the hand, $\mathbf{V}^h \in \mathbb{R}^N$ and those on the surface of the object, $\mathbf{V}^o \in \mathbb{R}^M$ for the *initial* frame. We can retrieve the hand vertices using the SMPL-X parameterization,

$$\mathbf{V}^n = SMPLX(\mathbf{P}^n). \tag{5}$$

For each subsequent frame, we then minimize the objective:

$$\mathbf{R}^*, \mathbf{T}^*, \mathbf{P}^{h*} = \min_{\mathbf{R}, \mathbf{T}, \mathbf{P}^{\mathbf{h}}} (\lambda_d E_d + \lambda_c E_c + \lambda_r E_r)$$
(6)

We use an energy term, E_d , to enforce the same inter-vertex distances between the hand and the object vertices in all the subsequent frames as in the first frame, as

$$E_d(\mathbf{R}, \mathbf{T}, \mathbf{P}^{\mathbf{h}}) = \left\| \operatorname{dist}(\mathbf{V}^{\mathbf{h}}, \mathbf{R}\mathbf{V}^{\mathbf{0}} + \mathbf{T}) - \mathbf{D} \right\|_2, \tag{7}$$

However, this term alone does not guarantee that the object is in contact with the hand in subsequent frames. This is mainly because in practice, the hand joints do not converge to plausible poses using

 E_d . We address this issue by introducing the contact term E_c , which forces the distance between the in-contact vertex-pairs of the first frame to be zero, as

$$E_c(\mathbf{P^h}) = \left\| \delta.\operatorname{dist}(\mathbf{V^h}, \mathbf{RV^o} + \mathbf{T}) \right\|_2.$$
(8)

Here, $\delta(\cdot, \cdot)$ is a contact indicator function for the elements of the distance matrix for which the distance is less than a threshold: $\delta(i, j) = 1$, if $\mathbf{D}_{i,j} < \tau$ and 0 otherwise, as we show in Fig. 4.

Finally, E_r consists of L2 regularizers to ensure that the object and hand poses do not deviate significantly from the previous frame and thus enforce temporal consistency, as

$$E_r(\mathbf{R}, \mathbf{T}, \mathbf{P}^{\mathbf{h}}) = \left\| \Delta \mathbf{R} + \Delta \mathbf{T} + \Delta \mathbf{P}^{\mathbf{h}} \right\|_2, \qquad (9)$$

where Δ signifies the difference in values between the current frame and the previous frame. We initialize the hand poses using a stateof-the-art grasp estimator proposed in [TCBT22]. The optimization routine course-corrects the initial estimates of the finger movements while also placing the object within the person's hands. Fig. 5 illustrates the optimization routine.

4. Implementation

This section describes our training and inference routines, and the implementation details for our generator network.

Training and Inference Routines. To maintain a fixed number of input frames for computational stability, to reduce the parameter load and associated training overheads, and to avoid overfitting to redundant frames, we represent our ground-truth motion sequences using T = 15 frames, taken at a sampling rate of 8-10 fps.

The encoders and the decoders inside our four modules use fullyconnected layers with skip connections, LeakyReLU activations and batch normalization [Aga18; BGSW18]. We use k = 4 past frames (optimized through experiments) to synthesize the next time steps. We train our CVAE-based Arm Synthesis and Body Synthesis Modules to minimize the KL divergence loss:

$$\mathcal{L}_{KL} = D_{KL} \left(\mathbf{q}_a \left(\mathbf{z}_{\mathbf{a}} | \mathbf{x}_{\mathbf{t}-\mathbf{k}:\mathbf{t}-1}, \phi \right) || \mathcal{N} \left(0, I \right) \right) + D_{KL} \left(\mathbf{q}_b \left(\mathbf{z}_{\mathbf{b}} | \mathbf{x}_{\mathbf{t}-\mathbf{k}:\mathbf{t}-1}, \phi \right) || \mathcal{N} \left(0, I \right) \right).$$
(10)

We compute the pose and the velocity reconstruction loss between the ground-truth rotations θ and the predicted rotations $\hat{\theta}$ as

$$\mathcal{L}_{rec} = \left\| \boldsymbol{\theta} - \hat{\boldsymbol{\theta}} \right\|_{1} + \left\| \Delta \boldsymbol{\theta} - \Delta \hat{\boldsymbol{\theta}} \right\|_{1}.$$
(11)

We train our model on the following weighted sum of these losses:

$$\mathcal{L} = \lambda_{KL} \mathcal{L}_{KL} + \lambda_p \mathcal{L}_{rec}, \qquad (12)$$

where λ_{KL} and λ_p are the weight parameters. We can then use the regressed body motion parameters $\hat{\mathbf{p}}$ to optimize the 6-DoF object positions at every time step.

At inference time, we synthesize motions for novel intent-object pairs and novel body shape parameters. We input an initial body pose, a 3D object placed within reach of the character, and an intended action to be performed with the object, and autoregressively synthesize the intent-based, full-body motion sequence.



Figure 5: **Object Position Optimization**. We optimize for the 6-DoF pose of the object such that it plausibly fits within the hands of the virtual character. We show three snapshots of such fitting after the 0^{th} , 500^{th} and the 1200^{th} iteration of our optimization.

Implementation Details. We train our model for 1500 epochs using the Adam Optimizer [KB14] with a base learning rate of 5×10^{-4} , and a batch size of 64, which takes roughly four hours on an NVIDIA A100-PCIE-40GB GPU. We decay the learning rate (LR) using a Reduce-on-plateau LR scheduler with a patience of 3 epochs and a decay rate of 0.999. We set $\lambda_{KL} = 0.001$, $\lambda_p = \lambda_d = 1.0$ and $\lambda_c = \lambda_r = 0.005$. During inference, synthesizing the full-body poses and the corresponding object positions for a motion sequence of 15 frames take approximately 1-1.5 minutes. Finally, we perform a linear interpolation on our generated frames to up-sample the motion to 30 frames per sequence for cleaner visualization. We have implemented our network, training, and inference using the PyTorch framework [PGC*17].

5. Experiments and Results

This section reports results of our experimental evaluation, including the dataset and the evaluation metrics we use and our ablation studies. Since there are no existing methods for generating fullbody human-object interactions, we use existing methods that generate full-body poses based only on action labels as our baselines.

5.1. Dataset

We use the GRAB dataset [TGBT20] consisting of whole-body grasping sequences performed by ten different subjects. The subjects interact with 51 different objects via four basic intents namely "use", "pass", "lift", and "off-hand". "Use" further has a subcategory of 26 different actions depicting plausible intent-object interactions such as drinking or pouring from a cup to taking picture with or browsing a camera. Following the split of [DFD22], we take subject 'S1' for validation, 'S10' for testing, and the remaining subjects 'S2' through 'S9' for training. The train, validation and test splits resepectively consist of 992, 197 and 136 sequences.

5.2. Baselines

We compare our results with ACTOR [PBV21], Action2Motion [GZW*20] and TEMOS [PBV22]. Since these methods have originally not been trained on the GRAB dataset, we retrain all of them for our setting. We re-train ACTOR and the Action2Motion methods for 1500 epochs (the same number of epochs we train our own model for, see Sec. 4) conditioned only on the action labels with no object information. For comparison with TEMOS, we create sentences of the form "A person <action> the <object>"(e.g., "a person eats the apple") to use as input sentences, and re-train the TEMOS model for 1500 epochs as well. For all the three motion synthesis methods, we apply our Object Optimizer Module to also generate the object positions for visual comparison.

5.3. Evaluation Metrics

We evaluate our method using the Mean Per Joint Positional Error (MPJPE). It measures the mean joint error over all time steps T as

$$MPJPE = \frac{1}{NT} \sum_{n \in N} \sum_{t \in T} \left\| J_t - \hat{J}_t \right\|_2, \qquad (13)$$

where \hat{J} are the joint positions computed from the synthesized SMPL-X parameters and N is the size of our test set. However, our task requires that the models synthesize a diverse set of plausible motions for any type of intent. Therefore, only calculating the Euclidean error with the ground-truth motion, does not provide a complete picture of their synthesis quality.

Therefore, to understand the overall motion distribution statistics, we use the Average Variance Error [GCO*21]. The Average Variance Error (AVE) computes the L_2 error between the variance of the joint positions and that of the ground truth as

AVE =
$$\frac{1}{N} \sum_{n \in N} \|\boldsymbol{\sigma} - \hat{\boldsymbol{\sigma}}\|_2$$
, with $\boldsymbol{\sigma} = \frac{1}{T-1} \sum_{t \in T} \left(J_t - \hat{J} \right)^2$, (14)

where \overline{J} is the mean pose over T time steps, σ is the ground-truth variance, and $\hat{\sigma}$ is the variance of the synthesized sequence.

We further evaluate the naturalness and the overall diversity of our generated motions using the Frechet Inception Distance (FID) [HRU*17], recognition accuracy, diversity and multimodality as done in ACTOR [PBV21] and Action2Motion [GZW*20]. We train a standard RNN action recognition classifier on the GRAB dataset, and use the final layer of the classifier as the motion feature extractor for calculating FID, diversity and multimodality.

Recognition accuracy indicates the correlation of the generated motions with their action types. Through diversity we measure variation in the generated motion's features across all action categories, whereas multimodality measures how generated motion's features diversify within each action type.

5.4. Ablation Studies

We compare the performance of our model with the following ablated versions:

- Ablation 1: One-hot vector representation for the input action labels: To study the effect of how the CLIP model influences the conditioning of the synthesized motion, we conduct an ablation where we train our Condition Encoder with the one-hot vectors of the input action label.
- Ablation 2: Training the Body Synthesis module without using the self-attention mapping. In this ablation, we exclude our position-encoded multi-head self-attention from the input of the Body Synthesis module of our framework to see how it influences the quality of our motion.
- Ablation 3: Training the whole-body instead of decoupling to the Arm Synthesis and the Body Synthesis Modules. We train

© 2023 The Author(s) Computer Graphics Forum © 2023 The Eurographics Association and John Wiley & Sons Ltd.



Figure 6: **Perceptual Study Evaluation.** We conduct a user study where participants answer two questions: "Which animation looks more realistic?" and "which animation best corresponds with the input instruction label?". We show them 30 randomly sampled motion sequences synthesized by our method and the two baselines, ACTOR [PBV21] and Action2Motion [GZW*20]. We see our method is chosen more than 80% times compared to ACTOR and Action2Motion.

the whole body movements in one module instead of separately synthesizing the arms and the rest of the body.

5.5. Quantitative Evaluation

Table 2 shows the FID, recognition accuracy, diversity, multimodality, MPJPE and AVE values on our test set compared with the three state-of-the-art methods of ACTOR [PBV21], Action2Motion [GZW*20], and TEMOS [PBV22]. We also include the ablated versions of our methods (Sec. 5.4) in our evaluation. We repeat each experiment 20 times as done in Action2Motion [GZW*20], and report a statistical interval with 95% confidence. We observe that our method shows significant improvements in each of the metrics compared to the existing methods and the ablated versions.

5.6. Perceptual Study

To evaluate the visual quality of our motions, we conduct a perceptual study where we compare our results with AC-TOR [PBV21] and Action2Motion [GZW*20]. With the exception of TEMOS [PBV22], which would quickly settle on the mean pose, the other two methods generated plausible full-body motions after retraining. We, therefore, exclude TEMOS from the user study. We conduct our perceptual study in the following two sections.

Comparison with Motion Synthesis Methods. In the first section, we displayed our results and the results from ACTOR and Action2Motion side-by-side in a random order along with the input instruction label. We asked the participants to answer the following two questions for each sequence: "Which motion looks the most realistic?" and "Which motion best corresponds with the input instruction label?". We collected answers for 30 such sequences from 75 participants. Fig. 6 illustrates the results of the study. In 80.7% responses, participants marked our method as the most realistic compared to ACTOR and Action2Motion. Likewise, 81.6% participants chose our method to have the best semantic fidelity with the instruction label. Upon examining the cases for which

A. Ghosh, R. Dabral, V. Golyanik, C. Theobalt, P. Slusallek / IMoS: Intent-Driven Full-Body Motion Synthesisfor Human-Object Interactions

Method	MPJPE (\downarrow)	AVE (\downarrow)	FID (\downarrow)	Accuracy (†)	Diversity (\rightarrow)	Multimodality (\rightarrow)
Real Motions (GT)	-	-	-	0.999	1.15 ± 0.015	0.30 ± 0.01
ACTOR	0.093	8.053	0.677	0.781	1.06 ± 0.015	0.19 ± 0.01
Action2Motion	0.113	8.266	1.087	0.694	1.10 ± 0.01	0.22 ± 0.01
TEMOS	0.103	9.981	1.211	0.228	0.83 ± 0.01	0.09 ± 0.01
Ablation 1	0.042	4.411	0.329	0.830	1.06 ± 0.015	0.21 ± 0.01
Ablation 2	0.042	4.772	0.387	0.824	1.10 ± 0.02	0.24 ± 0.02
Ablation 3	0.056	5.413	0.419	0.823	1.08 ± 0.01	0.25 ± 0.01
Ours	0.033	3.821	0.276	0.877	$\textbf{1.11} \pm 0.015$	$\textbf{0.28} \pm 0.015$

Table 2: **Quantitative Evaluation.** We compare with other motion synthesis methods, namely ACTOR [PBV21], Action2Motion [GZW*20] and TEMOS [PBV22], and three ablated versions of our model (Sec. 5.4). We evaluate the methods on the FID, recognition accuracy, diversity, multimodality, MPJPE and AVE metrics. \downarrow means lower values are better, \uparrow means higher values are better and \rightarrow means motions are better when the metric is closer to the ground-truth values.



Figure 7: Qualitative Results Showing Diversity in the Synthesized Motions. The two rows depict two diverse motion sequences generated by our model. We note that our method is able to generate different variations for the same instructions using either one or both hands along with plausible coordination of the head and the body. Please refer to the supplementary video for more results.

the participants preferred ACTOR instead of us, we found that it performed better for few actions like "screwing" the light bulb or "toasting" with the wineglass, where the motion does not need to have hand-to-eye or hand-to-mouth coordination. These actions do not include significant variations within the dataset and are, therefore, easy to overfit.

Comparison with Ground-Truth. While ACTOR and Action2Motion are methodically the closest to our approach, they were not originally designed to be conditioned according to our intent-based motion synthesis paradigm.

Therefore, to add an additional perspective on to the performance of our approach, we also asked the participants to compare our *best* synthesis results with the ground-truth in the second section. To establish an upper-bound on our performance, we chose our 10 best samples from various intent-object pairings to compare with the ground-truth. Again, we displayed our motions and the ground truth side-by-side in a random order. This time, we kept an extra option: "*cannot distinguish*". We observed that although our method is, expectedly, less preferred than the ground-truth motion (15.6% vs 36.9%), 47.5% of the responses rate our best syntheses as *indistinguishable* from the ground-truth in terms of realism. We also notice that participants rated our method to be more realistic than the ground-truth when it involves actions such as "eating" an apple with one hand, which has abundant training samples.

On the other hand, our method encounters difficulties when synthesizing intents involving high frequency wrist or finger movements such as "shaking" or "squeezing". We can attribute this to the fact that our L1 loss function (Eqn. (11)) tends to smooth out the high frequency components from the motion sequence and the

GRAB dataset does not have sufficient samples of these actions to train them separately.

5.7. Qualitative Evaluation

We show full qualitative results in our supplementary video. When qualitatively compared with the ablated versions (Sec. 5.4), we find that ablation 1 (one-hot vector instead of CLIP) and ablation 3 (training one module for whole-body) fail to synthesize precise hand-mouth or hand-eye coordination for actions such as "drinking" and "eating". Ablation 2 (without using self-attention mapping) lacks subtle body movements such as tilting back the head or bending the knee to pick up an object, which otherwise adds to plausibility of the motion. We further analyse our generated motions under the following headings:

Diversity Analysis. As we discussed initially (Sec. 1), generating diverse motion sequences for the same input instruction label is crucial for an immersive user experience. Fig 7 shows our result for two different sequences (left and right). Sampling from the variational latent space allows us to synthesize diverse motion sequences. In fig. 7 we show two different sequences: "taking picture" with a camera (left) and "eating" an apple (right). We show two variations of the same motions (upper and lower rows). We note that the variations are diverse in terms of the head, arms and torso movements, especially in the way they are angled to use the object. Our method benefits from the fact that we operate in the fullbody space and produces more natural results compared to naïvely performing a fixed mapping from the global hand pose parameters to the end effectors of the remaining body.

Synthesis of Both-Handed Interactions. Our method is the first to accurately synthesize full-body motions for both-handed interactions. We achieve this by decoupling the arm synthesis from full-body synthesis in our generator design (Sec. 3.2. The wrist and the elbow joints play a crucial role for tasks such as picking up an object with both hands or holding the object precisely. Learning the arm motions in a separate latent space helps our generator focus more on such precise synthesis.

Object Position Predictions for Off-handing Interactions. In addition to both-handed interactions, we encounter sequences in the GRAB dataset where the character passes an object from one hand to the other. It is non-trivial to optimize for the accurate object positions when the object is switching hands. For these sequences, we first compute the most-likely frame at which the switching takes place, and then transfer the optimized hand parameters to the other hand. Fig: 9 shows two such off-handing interactions with two different objects.

Plausibility of Head Motions. Similar to the motion of the fingers and the arms, the coordinated movement of the head and the hands also determines the synthesis quality. While recent works like GOAL [TCBT22] explicitly account for the head direction vector during network training and optimization, we observe that our model learns visually plausible head orientations and hand-head coordination without any explicit supervision. This raises the question whether explicit supervision is indeed necessary.

© 2023 The Author(s) Computer Graphics Forum © 2023 The Eurographics Association and John Wiley & Sons Ltd.



Figure 8: Examples of Imprecise Contacts in the GRAB Dataset [TGBT20]. We show five (ground-truth) frames where the body and the object are in contact. However, these contacts are not precise. The fingers do not touch the object for grasping the mug, the camera and the cup. For the cube and the toothpaste, we see inter-penetration of the hand with the object.

Figure 9: **Off-Handing.** We show two interactions of "offhanding" where the character passes the object from one hand to the other. Such interactions pose a unique optimization challenge when the object is switching hands.

6. Discussion and Limitations

Through quantitative evaluations and a perceptual study, we establish that our method synthesizes plausible motions corresponding to intended actions by a virtual character with a given object. While we can synthesize such motions for a variety of intents and objects, we observe certain failure cases for intents which are very rare and have high frequency in the wrist motion, *e.g.*, "squeeze", "shake" (see supplementary video). Additionally, our Object Optimizer Module (Sec. 3.2.4) optimizes the fingers and the object positions based on an initial distance between them. This assumption

works well with the intents in the GRAB dataset as most of them involve static grasps. However, dynamic grasping that involves hand slipping and relative motion between the object and the hands like "rotating" a cube or "stretching" an elastic band will be limited in our setting. We also note that the contacts between the body and the objects for all samples in the GRAB dataset are not fully precise, possibly due to the sparse marker-based motion-capture. In many sequences, we find that the fingers do not touch the object while grasping, or have inter-penetrations between the hand and the object (Fig. 8). Lastly, we do not address long term motion synthesis (in the order of minutes) involving a series of sequential actions performed with a given object.

7. Ethical Considerations

Our method does not support texture and fine appearance details and cannot be used to produce deceptive content. Our results are not photo-realistic by design and cannot be confused with real scenes. However, combining our technique with a method supporting more realistic texture could potentially raise ethical concerns in future.

8. Conclusion and Future Work

We presented the first full-body motion synthesis method for character-object interactions. Such a motion synthesis pipeline can become a useful practical tool in applications requiring large-scale character animations.

We demonstrate that a decoupling approach that separately models the arms and the body motions using CVAEs leads to measurable perceptual improvements as well as advances the state-of-theart on multiple quantitative evaluations. Our method places a special focus on synthesizing interactions involving *both* the hands, which also includes sequences where the object is exchanged between the hands.

In the future, we intend to extend our model to synthesize dynamic grasps along with full-body poses, such that the virtual character can change the grasp within a sequence. We also plan to explore more descriptive sentence embeddings for the interactions (*e.g.*, "a person passes the bowl using the right hand") in order to generate more precise and controllable motions.

References

- [Aga18] AGARAP, ABIEN FRED. "Deep learning using rectified linear units (relu)". arXiv preprint arXiv:1803.08375 (2018) 6.
- [AGK18] ANTOTSIOU, DAFNI, GARCIA-HERNANDO, GUILLERMO, and KIM, TAE-KYUN. "Task-oriented hand motion retargeting for dexterous manipulation imitation". Proceedings of the European Conference on Computer Vision (ECCV) Workshops. 2018 3.
- [AHC*18] AHN, H., HA, T., CHOI, Y., et al. "Text2Action: Generative Adversarial Synthesis from Language to Action". 2018 IEEE International Conference on Robotics and Automation (ICRA). 2018, 5915– 5920 2.
- [ALNM20] AHUJA, CHAITANYA, LEE, DONG WON, NAKANO, YUKIKO I, and MORENCY, LOUIS-PHILIPPE. "Style transfer for co-speech gesture animation: A multi-speaker conditional-mixture approach". *European Conference on Computer Vision*. Springer. 2020, 248–265 2.

- [AM19] AHUJA, C. and MORENCY, L. "Language2Pose: Natural Language Grounded Pose Forecasting". 2019 International Conference on 3D Vision (3DV). 2019, 719–728 2.
- [BCRM21] BHATTACHARYA, UTTARAN, CHILDS, ELIZA-BETH, REWKOWSKI, NICHOLAS, and MANOCHA, DINESH. "Speech2AffectiveGestures: Synthesizing Co-Speech Gestures with Generative Adversarial Affective Expression Learning". *Proceedings of the 29th ACM International Conference on Multimedia*. MM '21. New York, NY, USA: Association for Computing Machinery, 2021 2.
- [BGSW18] BJORCK, NILS, GOMES, CARLA P, SELMAN, BART, and WEINBERGER, KILIAN Q. "Understanding Batch Normalization". Advances in Neural Information Processing Systems. Ed. by BENGIO, S., WALLACH, H., LAROCHELLE, H., et al. 2018 6.
- [BHHF19] BRAHMBHATT, SAMARTH, HANDA, ANKUR, HAYS, JAMES, and FOX, DIETER. "Contactgrasp: Functional multi-finger grasp synthesis from contact". 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE. 2019, 2386–2393 3.
- [BI05] BORST, CHRISTOPH W and INDUGULA, ARUN P. "Realistic virtual grasping". *IEEE Proceedings. VR 2005. Virtual Reality*, 2005. IEEE. 2005, 91–98 3.
- [BRB*21] BHATTACHARYA, UTTARAN, REWKOWSKI, NICHOLAS, BANERJEE, ABHISHEK, et al. "Text2Gestures: A Transformer-Based Network for Generating Emotive Body Gestures for Virtual Agents". 2021 IEEE Conference on Virtual Reality and 3D User Interfaces (IEEE VR). IEEE, 2021 2.
- [BTT*20] BRAHMBHATT, SAMARTH, TANG, CHENGCHENG, TWIGG, CHRISTOPHER D., et al. "ContactPose: A Dataset of Grasps with Object Contact and Hand Pose". *The European Conference on Computer Vision (ECCV)*. Aug. 2020 3.
- [BXP*22] BHATNAGAR, BHARAT LAL, XIE, XIANGHUI, PETROV, ILYA, et al. "BEHAVE: Dataset and Method for Tracking Human Object Interactions". *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR). 2022 3.
- [CKA*22] CHRISTEN, SAMMY, KOCABAS, MUHAMMED, AKSAN, EMRE, et al. "D-Grasp: Physically Plausible Dynamic Grasp Synthesis for Hand-Object Interactions". Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 2022 2, 3.
- [DFD22] DILLER, CHRISTIAN, FUNKHOUSER, THOMAS, and DAI, AN-GELA. "Forecasting Characteristic 3D Poses of Human Actions". (2022) 3, 6.
- [DKB*10] DETRY, RENAUD, KRAFT, DIRK, BUCH, ANDERS GLENT, et al. "Refining grasp affordance models by experience". 2010 IEEE International Conference on Robotics and Automation. IEEE. 2010, 2287– 2293 3.
- [DSJ*21] DABRAL, RISHABH, SHIMADA, SOSHI, JAIN, ARJUN, et al. "Gravity-Aware Monocular 3D Human-Object Reconstruction". International Conference on Computer Vision (ICCV). 2021 3.
- [ES03] ELKOURA, GEORGE and SINGH, KARAN. "Handrix: animating the human hand". Proceedings of the 2003 ACM SIG-GRAPH/Eurographics symposium on Computer animation. 2003 3.
- [GCO*21] GHOSH, ANINDITA, CHEEMA, NOSHABA, OGUZ, CENNET, et al. "Synthesis of Compositional Animations From Textual Descriptions". Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). Oct. 2021, 1396–1406 2, 7.
- [GPM*20] GOODFELLOW, IAN, POUGET-ABADIE, JEAN, MIRZA, MEHDI, et al. "Generative adversarial networks". *Communications of* the ACM 63.11 (2020), 139–144 3.
- [GTT*21] GRADY, PATRICK, TANG, CHENGCHENG, TWIGG, CHRISTO-PHER D., et al. "ContactOpt: Optimizing Contact To Improve Grasps". Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). June 2021, 1471–1481 3.
- [GZW*20] GUO, CHUAN, ZUO, XINXIN, WANG, SEN, et al. "Action2motion: Conditioned generation of 3d human motions". Proceedings of the 28th ACM International Conference on Multimedia. 2020 2, 3, 6–8.

© 2023 The Author(s) Computer Graphics Forum © 2023 The Eurographics Association and John Wiley & Sons Ltd.

- [GZZ*22] GUO, CHUAN, ZOU, SHIHAO, ZUO, XINXIN, et al. "Generating Diverse and Natural 3D Human Motions From Text". Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). June 2022, 5152–5161 2, 3.
- [HCTB19] HASSAN, MOHAMED, CHOUTAS, VASILEIOS, TZIONAS, DIMITRIOS, and BLACK, MICHAEL J. "Resolving 3D Human Pose Ambiguities with 3D Scene Constraints". *International Conference on Computer Vision*. 2019 3.
- [HCV*21] HASSAN, MOHAMED, CEYLAN, DUYGU, VILLEGAS, RUBEN, et al. "Stochastic Scene-Aware Motion Prediction". Proceedings of the International Conference on Computer Vision 2021. 2021 3.
- [HL06] HSIAO, KAIJEN and LOZANO-PEREZ, TOMAS. "Imitation learning of whole-body grasps". 2006 IEEE/RSJ international conference on intelligent robots and systems. IEEE. 2006, 5657–5662 3.
- [HMLC09] HANSER, EVA, MC KEVITT, PAUL, LUNNEY, TOM, and CONDELL, JOAN. "Scenemaker: Intelligent multimodal visualisation of natural language scripts". *Irish Conference on Artificial Intelligence and Cognitive Science*. Springer. 2009, 144–153 2.
- [HRU*17] HEUSEL, MARTIN, RAMSAUER, HUBERT, UNTERTHINER, THOMAS, et al. "Gans trained by a two time-scale update rule converge to a local nash equilibrium". *Advances in neural information processing systems* 30 (2017) 7.
- [HVT*19] HASSON, YANA, VAROL, GÜL, TZIONAS, DIMITRIOS, et al. "Learning joint reconstruction of hands and manipulated objects". *CVPR*. 2019 3.
- [JLWW21] JIANG, HANWEN, LIU, SHAOWEI, WANG, JIASHUN, and WANG, XIAOLONG. "Hand-object contact consistency reasoning for human grasps generation". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2021, 11107–11116 2, 3.
- [KB14] KINGMA, DIEDERIK P and BA, JIMMY. "Adam: A method for stochastic optimization". arXiv preprint arXiv:1412.6980 (2014) 6.
- [KDCI10] KRUG, ROBERT, DIMITROV, DIMITAR, CHARUSTA, KRZYSZTOF, and ILIEV, BOYKO. "On the efficient computation of independent contact regions for force closure grasps". 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE. 2010, 586–591 3.
- [KP15] KIM, JUN-SIK and PARK, JUNG-MIN. "Physics-based hand interaction with virtual objects". 2015 IEEE International Conference on Robotics and Automation (ICRA). IEEE. 2015, 3814–3819 3.
- [KTS*21] KWON, TAEIN, TEKIN, BUGRA, STÜHMER, JAN, et al. "H2O: Two Hands Manipulating Objects for First Person Interaction Recognition". Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). 2021 3.
- [KW14] KINGMA, DIEDERIK and WELLING, MAX. "Auto-Encoding Variational Bayes". Dec. 2014 3.
- [KYZ*20] KARUNRATANAKUL, KORRAWE, YANG, JINLONG, ZHANG, YAN, et al. "Grasping Field: Learning Implicit Representations for Human Grasps". 8th International Conference on 3D Vision. IEEE, Nov. 2020, 333–344 2, 3.
- [LFP07] LI, YING, FU, JIAXIN L, and POLLARD, NANCY S. "Data-driven grasp synthesis using shape matching and task-based pruning". *IEEE Transactions on visualization and computer graphics* 13.4 (2007), 732–747 3.
- [LLW*21] LIU, ZHENGUANG, LYU, KEDI, WU, SHUANG, et al. "Aggregated multi-gans for controlled 3d human motion prediction". Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 35. 3. 2021, 2225–2232 3.
- [LPX*19] LIU, MIN, PAN, ZHERONG, XU, KAI, et al. "Generating grasp poses for a high-dof gripper using neural networks". 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE. 2019, 1518–1525 3.
- [LRX*21] LI, WEIPENG, REN, BOYUAN, XU, HAOYUE, et al. "Auto-Dance: Music Driven Dance Generation". May 2021, 55–59 2.

© 2023 The Author(s)

Computer Graphics Forum © 2023 The Eurographics Association and John Wiley & Sons Ltd.

- [LWC*18] LIN, ANGELA S, WU, LEMENG, CORONA, RODOLFO, et al. "Generating animated videos of human activities from natural language descriptions". (2018) 2.
- [LWM21] LIN, FANQING, WILHELM, CONNOR, and MARTINEZ, TONY. "Two-Hand Global 3D Pose Estimation Using Monocular RGB". Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). Jan. 2021, 2373–2381 3.
- [LYC*20] LI, JIAMAN, YIN, YIHANG, CHU, HANG, et al. "Learning to generate diverse dance motions with transformer". *arXiv preprint arXiv:2008.08171* (2020) 2.
- [LYL*19] LEE, HSIN-YING, YANG, XIAODONG, LIU, MING-YU, et al. "Dancing to Music". Advances in Neural Information Processing Systems. Ed. by WALLACH, H., LAROCHELLE, H., BEYGELZIMER, A., et al. Vol. 32. Curran Associates, Inc., 2019 2.
- [LZCvdP20] LING, HUNG YU, ZINNO, FABIO, CHENG, GEORGE, and van de PANNE, MICHIEL. "Character Controllers Using Motion VAEs". ACM Trans. Graph. (2020) 2, 3.
- [MBR17] MARTINEZ, JULIETA, BLACK, MICHAEL J, and ROMERO, JAVIER. "On human motion prediction using recurrent neural networks". *Proceedings of the IEEE conference on computer vision and pattern* recognition. 2017, 2891–2900 3.
- [PBV21] PETROVICH, MATHIS, BLACK, MICHAEL J., and VAROL, GÜL. "Action-Conditioned 3D Human Motion Synthesis with Transformer VAE". International Conference on Computer Vision (ICCV). Oct. 2021, 10985–10995 3, 6–8.
- [PBV22] PETROVICH, MATHIS, BLACK, MICHAEL J., and VAROL, GÜL. "TEMOS: Generating diverse human motions from textual descriptions". arXiv (Apr. 2022) 2, 3, 6–8.
- [PCG*19] PAVLAKOS, GEORGIOS, CHOUTAS, VASILEIOS, GHORBANI, NIMA, et al. "Expressive Body Capture: 3D Hands, Face, and Body from a Single Image". Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). June 2019, 10975–10985 4.
- [PG16] PINTO, LERREL and GUPTA, ABHINAV. "Supersizing selfsupervision: Learning to grasp from 50k tries and 700 robot hours". 2016 IEEE international conference on robotics and automation (ICRA). IEEE. 2016, 3406–3413 3.
- [PGA18] PAVLLO, DARIO, GRANGIER, DAVID, and AULI, MICHAEL. "Quaternet: A quaternion-based recurrent model for human motion". arXiv preprint arXiv:1805.06485 (2018) 3.
- [PGC*17] PASZKE, ADAM, GROSS, SAM, CHINTALA, SOUMITH, et al. "Automatic Differentiation in PyTorch". *NeurIPS 2017 Workshop on Autodiff.* 2017 6.
- [PRB*18] PUIG, XAVIER, RA, KEVIN, BOBEN, MARKO, et al. "VirtualHome: Simulating Household Activities via Programs". Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). June 2018 2.
- [RA15] REDMON, JOSEPH and ANGELOVA, ANELIA. "Real-time grasp detection using convolutional neural networks". 2015 IEEE international conference on robotics and automation (ICRA). IEEE. 2015, 1316– 1322 3.
- [RBH*21] REMPE, DAVIS, BIRDAL, TOLGA, HERTZMANN, AARON, et al. "HuMoR: 3D Human Motion Model for Robust Pose Estimation". *International Conference on Computer Vision (ICCV)*, 2021 2, 3.
- [RKH*] RADFORD, ALEC, KIM, JONG WOOK, HALLACY, CHRIS, et al. "Learning Transferable Visual Models From Natural Language Supervision". Proceedings of the 38th International Conference on Machine Learning 2.
- [RTB17] ROMERO, JAVIER, TZIONAS, DIMITRIOS, and BLACK, MICHAEL J. "Embodied Hands: Modeling and Capturing Hands and Bodies Together". ACM Transactions on Graphics, (Proc. SIGGRAPH Asia). 245:1–245:17 36.6 (Nov. 2017), 245:1–245:17 3.
- [SKK12] SEO, JUNGWON, KIM, SOONKYUM, and KUMAR, VIJAY. "Planar, bimanual, whole-arm grasping". 2012 IEEE International Conference on Robotics and Automation. 2012, 3271–3277 3.

- [SLY15] SOHN, KIHYUK, LEE, HONGLAK, and YAN, XINCHEN. "Learning Structured Output Representation using Deep Conditional Generative Models". Advances in Neural Information Processing Systems. 2015 2.
- [SSR07] SUNG, KELVIN, SHIRLEY, PETER, and ROSENBERG, BECKY. "Experiencing aspects of games programming in an introductory computer graphics class". Mar. 2007, 249–253 2.
- [SZKS19] STARKE, SEBASTIAN, ZHANG, HE, KOMURA, TAKU, and SAITO, JUN. "Neural State Machine for Character-Scene Interactions". *ACM Trans. Graph.* (2019) 3.
- [TCBT22] TAHERI, OMID, CHOUTAS, VASSILEIOS, BLACK, MICHAEL J., and TZIONAS, DIMITRIOS. "GOAL: Generating 4D Whole-Body Motion for Hand-Object Grasping". *Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022 2, 3, 6, 9.
- [TGBT20] TAHERI, OMID, GHORBANI, NIMA, BLACK, MICHAEL J., and TZIONAS, DIMITRIOS. "GRAB: A Dataset of Whole-Body Human Grasping of Objects". *European Conference on Computer Vision* (ECCV). 2020 2, 3, 6, 9.
- [TGH*22] TEVET, GUY, GORDON, BRIAN, HERTZ, AMIR, et al. "MotionCLIP: Exposing Human Motion Generation to CLIP Space". arXiv preprint arXiv:2203.08063 (2022) 4.
- [VSP*17] VASWANI, ASHISH, SHAZEER, NOAM, PARMAR, NIKI, et al. "Attention is all you need". Advances in neural information processing systems 30 (2017) 3.
- [WWZ*22] WU, YAN, WANG, JIAHAO, ZHANG, YAN, et al. "SAGA: Stochastic Whole-Body Grasping with Contact". Proceedings of the European Conference on Computer Vision (ECCV). 2022 2–4.
- [XJMS21] XU, XIANG, JOO, HANBYUL, MORI, GREG, and SAVVA, MANOLIS. "D3D-HOI: Dynamic 3D Human-Object Interactions from Videos". arXiv preprint arXiv:2108.08420 (2021) 3.
- [YK20] YUAN, YE and KITANI, KRIS. "Dlow: Diversifying latent flows for diverse human motion prediction". *European Conference on Computer Vision*. Springer. 2020, 346–364 3.
- [ZBL*19] ZHOU, YI, BARNES, CONNELLY, LU, JINGWAN, et al. "On the Continuity of Rotation Representations in Neural Networks". Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). June 2019 4.
- [ZHT*21] ZHANG, XIONG, HUANG, HONGSHENG, TAN, JIANCHAO, et al. "Hand Image Understanding via Deep Multi-Task Learning". Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). 2021, 11281–11292 3.
- [ZLM*19] ZHANG, XIONG, LI, QIANG, MO, HONG, et al. "End-to-end hand mesh recovery from a monocular rgb image". Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). 2019, 2354–2364 3.
- [ZPJ*20] ZHANG, JASON Y., PEPOSE, SAM, JOO, HANBYUL, et al. "Perceiving 3D Human-Object Spatial Arrangements from a Single Image in the Wild". *European Conference on Computer Vision (ECCV)*. 2020 3.
- [ZYSK21] ZHANG, HE, YE, YUTING, SHIRATORI, TAKAAKI, and KO-MURA, TAKU. "ManipNet: Neural Manipulation Synthesis with a Hand-Object Spatial Representation". ACM Trans. Graph. (2021) 2, 3.