

MACS: Mass Conditioned 3D Hand and Object Motion Synthesis

Soshi Shimada^{1,2,*} Franziska Mueller³ Jan Bednarik³ Bardia Doosti³ Bernd Bickel³
Danhang Tang³ Vladislav Golyanik¹ Jonathan Taylor³ Christian Theobalt^{1,2} Thabo Beeler³

¹MPI for Informatics, SIC ²VIA Research Center ³Google

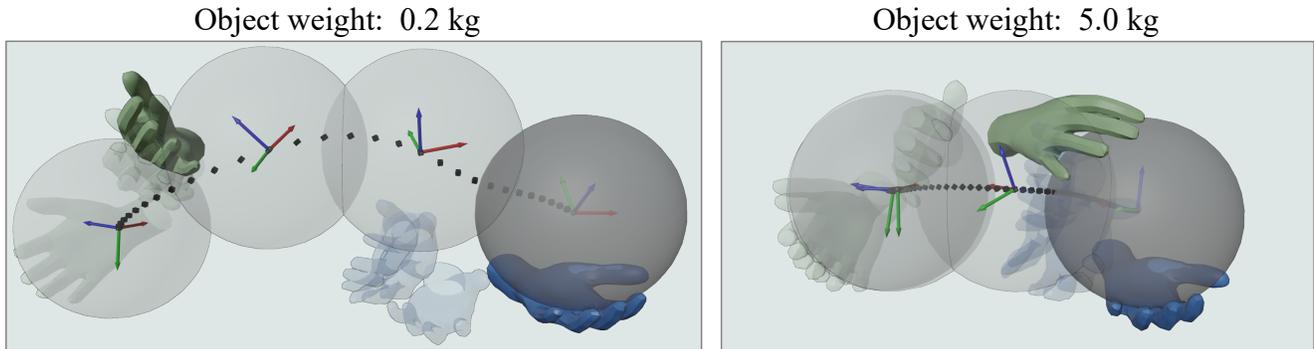


Figure 1. Example visualizations of 3D object manipulation synthesized by our method MACS. Conditioning object mass values of 0.2kg (left) and 5.0kg (right) are given to the model for the action type "passing from one hand to another". MACS plausibly reflects the mass value in the synthesized 3D motions.

Abstract

The physical properties of an object, such as mass, significantly affect how we manipulate it with our hands. Surprisingly, this aspect has so far been neglected in prior work on 3D motion synthesis. To improve the naturalness of the synthesized 3D hand-object motions, this work proposes MACS—the first MAAss Conditioned 3D hand and object motion Synthesis approach. Our approach is based on cascaded diffusion models and generates interactions that plausibly adjust based on the object’s mass and interaction type. MACS also accepts a manually drawn 3D object trajectory as input and synthesizes the natural 3D hand motions conditioned by the object’s mass. This flexibility enables MACS to be used for various downstream applications, such as generating synthetic training data for ML tasks, fast animation of hands for graphics workflows, and generating character interactions for computer games. We show experimentally that a small-scale dataset is sufficient for MACS to reasonably generalize across interpolated and extrapolated object masses unseen during the training. Furthermore, MACS shows moderate generalization to unseen objects, thanks to the mass-conditioned contact labels generated by our surface contact synthesis model ConNet. Our comprehensive user study confirms that the synthesized 3D hand-object interactions are highly plausible and realistic.

*Work done while at Google.

Project page link: <https://vcai.mpi-inf.mpg.de/projects/MACS/>

1. Introduction

Hand-object interaction plays an important role in our daily lives, involving the use of our hands in a variety of ways such as grasping, lifting, and throwing. It is crucial for graphics applications (e.g. AR/VR, avatar communication and character animation) to synthesize or capture physically plausible interactions for their enhanced realism. Therefore, there has been a growing interest in this field of research, and a significant amount of work has been proposed in grasp synthesis [11, 15, 18, 19, 31], object manipulation [4, 9, 22, 38, 41], 3D reconstruction [6, 14, 20, 23, 28, 33, 36], graph refinement [8, 24, 44] and contact prediction [3].

Because of the high-dimensionality of the hand models and inconsistent object shape and topology, synthesizing plausible 3D hand-object interaction is challenging. Furthermore, errors of even a few millimeters can cause collisions or floating-object artefacts that immediately convey an unnatural impression to the viewer. Some works tackle the static grasp synthesis task using an explicit hand model [11, 18, 31] or an implicit representation [15]. However, considering the static frame alone is not sufficient to integrate the method into real-world applications such as AR/VR as it lacks information of the inherent scene dynamics. Recently, several works have been proposed to syn-

thesize the hand and object interactions as a continuous sequence [4, 41, 44]. However, none of the state-of-the-art work explicitly considers an object’s mass when generating hand-object interactions. Real-life object manipulation, however, is substantially influenced by the mass of the objects we are interacting with. For example, we tend to grab light objects using our fingertips, whereas with heavy objects oftentimes the entire palm is in contact with the object. Manually creating such animations is tedious work requiring artistic skills. In this work, we propose *MACS*, *i.e.*, the first learning-based *mass conditioned* object manipulation synthesis method. The generated object manipulation naturally adopts its behavior depending on the object mass value. *MACS* can synthesize such mass conditioned interactions given a trajectory plus action label (*e.g.*, throw or move). The trajectory itself may also be generated conditioned on the action label and mass using the proposed cascaded diffusion model, or alternatively manually specified.

Specifically, given the action label and mass value as conditions, our cascaded diffusion model synthesizes the object trajectories as the first step. The synthesized object trajectory and mass value further condition a second diffusion model that synthesizes 3D hand motions and hand contact labels. After the final optimization step, *MACS* returns diverse and physically plausible object manipulation animations. We also demonstrate a simple but effective data capture set-up to produce a 3D object manipulation dataset with corresponding mass values. The contributions of our work are as follows:

- The first approach to synthesize *mass-conditioned* object manipulations in 3D. Our setting includes two hands and a single object of varying mass.
- A cascaded denoising diffusion model for generating trajectories of hands and objects allowing different types of conditioning inputs. Our approach can both synthesize new object trajectories and operate on user-provided trajectories (in this case, the object trajectory synthesis part is skipped).
- A new component for introducing plausible dynamics into user-provided trajectories.

Our experiments confirm that *MACS* synthesizes qualitatively and quantitatively more plausible 3D object manipulations compared with other baselines. *MACS* shows plausible manipulative interactions even for mass values vastly different from those seen during the training.

2. Related Work

There has been a significant amount of research in the field of 3D hand-object interaction motion synthesis. Here, we will review some of the most relevant works in this area. Grasp synthesis works are discussed in Sec. 2.1 and works that generate hand-object manipulation sequences in Sec. 2.2. Lastly, closely related recent diffusion model

based synthesis approaches are discussed in Sec. 2.3.

2.1. Grasp Synthesis

Synthesising physically plausible and natural grasps bears a lot of potential downstream applications. Thus, many works in this field have been proposed in computer graphics and vision [9, 19, 24, 38, 41], and robotics community [18, 35]. ContactOpt [11] utilizes a differentiable contact model to obtain a plausible grasp from a hand and object mesh. Karunratanakul *et al.* [15] proposed a *grasping field* for a grasp synthesis where hand and object surfaces are implicitly represented using a signed distance field. Zhou *et al.* [44] proposed a learning-based object grasp refinement method given noisy hand grasping poses. GOAL [32] synthesizes a whole human body motion with grasps along with plausible head directions. These works synthesize natural hand grasp on a variety of objects. However, unlike the methods in this class, we synthesize a sequential object manipulation, changing not only the hand pose but also object positions bearing plausible hand-object interactions.

2.2. Object Manipulation

Synthesising a sequence for object manipulation is challenging since the synthesized motions have to contain temporal consistency and plausible dynamics considering the continuous interactions. Ghosh *et al.* [9] proposed a human-object interaction synthesis algorithm associating the intentions and text inputs. ManipNet [41] predicts dexterous object manipulations with one/two hands given 6 DoF of hands and object trajectory from a motion tracker. CAMS [43] synthesizes hand articulations given a sequence of interacting object positions. Unlike these approaches, our algorithm **synthesizes** the 6 DoF of the hands and objects as well as the finger articulations affected by the conditioned mass values. D-Grasp [4] is a reinforcement learning-based method that leverages a physics simulation to synthesize a dynamic grasping motion that consists of approaching, grasping and moving a target object. In contrast to D-Grasp, our method consists of a cascaded diffusion model architecture and has explicit control over the object mass value that influences the synthesized interactions. Furthermore, D-Grasp uses a predetermined target grasp pose and therefore does not faithfully adjust its grasp based on the mass value in the simulator unlike ours.

2.3. Diffusion Model based Synthesis

Recently, diffusion model [29] based synthesis approaches have been receiving growing attention due to their promising results in a variety of research fields *e.g.* image generation tasks [13, 26, 27], audio synthesis [17], motion synthesis [7, 34, 40, 42] and 3D character generation from texts [25]. MDM [34] shows the 3D human motion synthesis and inpainting tasks from conditional action or text inputs

utilizing a transformer-based architecture allowing the integration of the geometric loss terms during the training. Our method is the first diffusion model based approach that synthesizes hand-object interactions. Furthermore, unlike the existing works in the literature, we condition the synthesized motions on a physical property, *i.e.*, object mass.

3. Method

Our goal is to synthesize 3D motion sequences of two hands interacting with an object whose mass affects both the trajectory of the object and the way the hands grasp it. The inputs of this method are a conditional scalar mass value and optionally a one-hot coded action label and/or a manually drawn object trajectory. Our method synthesizes a *motion* represented as N successive pairs of 3D hands and object poses. To this end, we employ denoising diffusion models (DDM) [29] for 3D hand motion and object trajectory synthesis; see Fig. 2 for the overview. We first describe our mathematical modeling and assumptions in Sec. 3.1. In Secs. 3.2 and 3.3, we provide details of our hand motion synthesis network *HandDiff* and trajectory synthesis algorithm *TrajDiff*, respectively. We describe the method to synthesize the 3D motions given user input trajectory in Sec. 3.3.2. The details of network architectures and training are elaborated in our supplementary material.

3.1. Assumptions, Modelling and Preliminaries

In this work, we assume that the target object is represented as a mesh. 3D hands are represented with a consistent topology, which is described in the following paragraph.

Hand and Object Modelling To represent 3D hands, we employ the hand model from GHUM [37] which is a non-linear parametric model learned from large-scale 3D human scans. The hand model from GHUM defines the 3D hand mesh as a differentiable function $\mathcal{M}(\tau, \phi, \theta, \beta)$ of global root translation $\tau \in \mathbb{R}^3$, global root orientation $\phi \in \mathbb{R}^6$ represented in 6D rotation representation [45], pose parameters $\theta \in \mathbb{R}^{90}$ and shape parameters $\beta \in \mathbb{R}^{16}$. We employ two GHUM hand models to represent left and right hands, which return hand vertices $\mathbf{v} \in \mathbb{R}^{3l}$ ($l = 1882 = 941 \cdot 2$) and 3D hand joints $\mathbf{j} \in \mathbb{R}^{3K}$ ($K = 42 = 21 \cdot 2$). The object pose is represented by its 3D translation $\tau_{\text{obj.}} \in \mathbb{R}^3$ and rotation $\phi_{\text{obj.}} \in \mathbb{R}^6$. Our method *MACS* synthesizes N successive (i) 3D hand motions represented by the hand vertices $\mathbf{V} = \{\mathbf{v}_1, \dots, \mathbf{v}_N\} \in \mathbb{R}^{N \times 3l}$ and hand joints $\mathbf{J} = \{\mathbf{j}_1, \dots, \mathbf{j}_N\} \in \mathbb{R}^{N \times 3K}$, and (ii) optionally object poses

$$\Phi = \{\Phi_1, \dots, \Phi_N\} \in \mathbb{R}^{N \times (3+6)}, \quad (1)$$

where $\Phi_i = [\tau_{\text{obj.},i}, \phi_{\text{obj.},i}]$. The object pose is defined in a fixed world frame \mathcal{F} , and the global hand translations are represented relative to the object center position. The global hand rotations are represented relative to \mathcal{F} .

Denoising Diffusion Model The recently proposed Denoising Diffusion Probabilistic Model (DDPM) [13] has shown compelling results both in image synthesis tasks and in motion generation tasks [34]. Compared to other existing generative models (*e.g.*, VAE [30] or GAN [10]) that are often employed for motion synthesis tasks, the training of DDPM is simple, as it is not subject to the notorious mode collapse while generating motions of high quality and diversity.

Following the formulation by Ho et al. [13], the forward diffusion process is defined as a Markov process adding Gaussian noise in each step. The noise injection is repeated T times. Next, let $\mathbf{X}^{(0)}$ be the original ground-truth (GT) data (without noise). Then, the forward diffusion process is defined by a distribution $q(\cdot)$:

$$q\left(\mathbf{X}^{(1:T)} \mid \mathbf{X}^{(0)}\right) = \prod_{t=1}^T q\left(\mathbf{X}^{(t)} \mid \mathbf{X}^{(t-1)}\right), \quad (2)$$

$$q\left(\mathbf{X}^{(t)} \mid \mathbf{X}^{(t-1)}\right) = \mathcal{N}\left(\mathbf{X}^{(t)} \mid \sqrt{1 - \beta_t} \mathbf{X}^{(t-1)}, \beta_t \mathbf{I}\right), \quad (3)$$

where β_t are constant hyperparameters (scalars) that are fixed per each diffusion time step t . Using a reparametrization technique, we can sample $\mathbf{X}^{(t)}$ using the original data $\mathbf{X}^{(0)}$ and standard Gaussian noise $\epsilon \sim \mathcal{N}(0, I)$:

$$\mathbf{X}^{(t)} = \sqrt{\alpha_t} \mathbf{X}^{(0)} + \sqrt{1 - \alpha_t} \epsilon, \quad (4)$$

where $\alpha_t = \prod_{i=1}^t (1 - \beta_i)$. The network is trained to reverse this process by denoising on each diffusion time step starting from a standard normal distribution $\mathbf{X}^{(T)} \sim \mathcal{N}(0, I)$:

$$p\left(\mathbf{X}^{(0:T)}\right) = p\left(\mathbf{X}^{(T)}\right) \prod_{t=1}^T p\left(\mathbf{X}^{(t-1)} \mid \mathbf{X}^{(t)}\right), \quad (5)$$

where $p\left(\mathbf{X}^{(t-1)} \mid \mathbf{X}^{(t)}\right)$ denotes the conditional probability distribution estimated from the network output. From Eq. (5), we obtain the meaningful generated result \mathbf{X}^* after T times of denoising process. that follows the data distribution of the training dataset.

In the formulation of DDPM [13], the network is trained to predict the added noises on the data for the reverse diffusion process. The *simple* loss term is formulated as

$$\mathcal{L}_{\text{simple}} = E_{\epsilon, t \sim [1, T]} \left[\left\| \epsilon - \epsilon_{\theta}\left(\mathbf{X}^{(t)}, t, c\right) \right\|_2^2 \right], \quad (6)$$

where c denotes an optional conditioning vector. The loss term of Eq. (6) drives the network ϵ_{θ} towards predicting the added noise. Training the network with Eq. (6) alone already generates highly diverse motions.

In our case \mathbf{X}^* represents sequences of 3D points corresponding to the synthesized motion trajectories (for hands

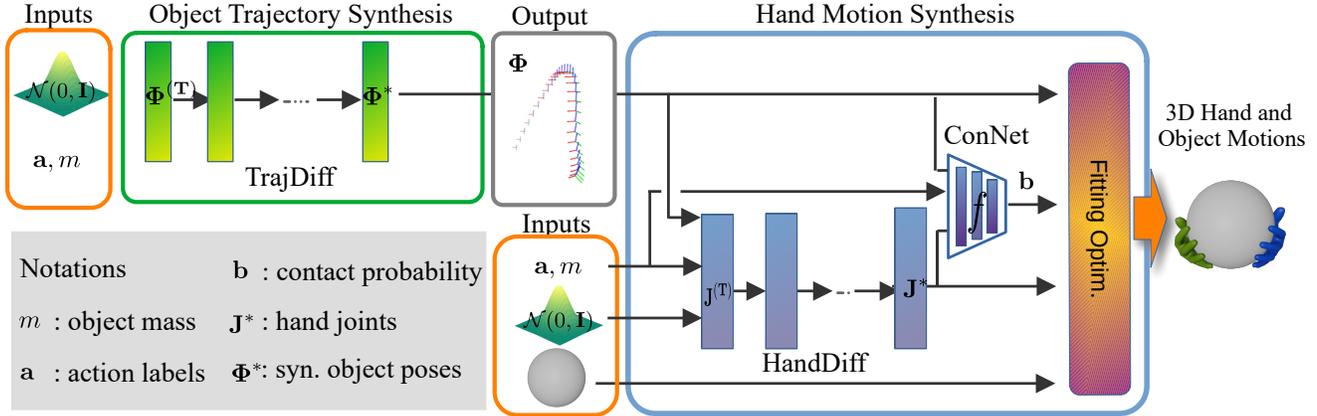


Figure 2. **The proposed framework.** The object trajectory synthesis stage accepts as input the conditional mass value m and action label \mathbf{a} along with a Gaussian noise sampled from $\mathcal{N}(0, \mathbf{I})$, and outputs an object trajectory. The hand motion synthesis stage accepts \mathbf{a}, m and the synthesized trajectory as conditions along with a gaussian noise sampled from $\mathcal{N}(0, \mathbf{I})$. ConNet in this stage estimates the per-vertex hand contacts from the synthesized hand joints, object trajectory and conditioning values \mathbf{a}, m . The final fitting optimization step returns a set of 3D hand meshes that plausibly interact with the target object.

and objects). Unfortunately, Eq. (6) alone often leads to artifacts in the generated sequences such as joint jitters and varying bone length when applied to motion synthesis. To improve the plausibility of the generated results, Dabral et al. [7] proposed an algorithm to integrate the explicit geometric loss terms into the training of DDPM. At an arbitrary diffusion time step t , we can obtain the approximated original data $\hat{\mathbf{X}}^{(0)}$ using the estimated noise from ϵ_θ instead of ϵ in Eq. (4) and solving for $\hat{\mathbf{X}}^{(0)}$:

$$\hat{\mathbf{X}}^{(0)} = \frac{1}{\sqrt{\alpha}} \mathbf{X}^{(t)} - \left(\sqrt{\frac{1}{\alpha} - 1} \right) \epsilon_\theta \left(\mathbf{X}^{(t)}, t, c \right). \quad (7)$$

During the training, geometric penalties can be applied on $\hat{\mathbf{X}}^{(0)}$ so as to prevent the aforementioned artifacts. In the following sections, we follow the mathematical notations of DDPM literature [7, 13] as much as possible. The approximated set of hand joints and object poses obtained from Eq. (7) are denoted $\hat{\mathbf{J}}^{(0)}$ and $\hat{\Phi}^{(0)}$, respectively. Similarly, the synthesized set of meaningful hand joints and object poses obtained from the reverse diffusion process Eq. (5) are denoted \mathbf{J}^* and Φ^* , respectively.

3.2. Hand 3D Motion Synthesis

Our DDPM-based architectures *HandDiff* $\mathcal{H}(\cdot)$ and *TrajDiff* $\mathcal{T}(\cdot)$ are based on the stable diffusion architecture [26] with simple 1D and 2D convolution layers (see our supplementary for more details). During the training, we follow the formulation of Dabral et al. [7] described in Sec. 3.1 to introduce geometric penalties on $\hat{\mathbf{J}}^{(0)} \in \mathbb{R}^{N \times 3K}$ and $\hat{\Phi}^{(0)} \in \mathbb{R}^{N \times 9}$ combined with the simple loss described in Eq. (6).

Hand Keypoints Synthesis In this stage, we synthesize a set of 3D hand joints and per-vertex hand contact probabilities. Knowing the contact positions on hands substantially helps to reduce the implausible “floating object” artifacts of the object manipulation (see Sec.4 for the ablations). The synthesized 3D hand joints and contact information are further sent to the final fitting optimization stage where we obtain the final hand meshes considering the plausible interactions between the hands and the object.

Our diffusion model based *HandDiff* $\mathcal{H}(\cdot)$ accepts as inputs a 3D trajectory $\Phi \in \mathbb{R}^{N \times (3+6)}$ and mass scalar value m where N is the number of frames of the sequence. From the reverse diffusion process of $\mathcal{H}(\cdot)$, we obtain the synthesized set of 3D joints $\mathbf{J}^* \in \mathbb{R}^{N \times 3K}$. Φ can be either synthesized by *TrajDiff* $\mathcal{T}(\cdot)$ (Sec. 3.3.1) or manually provided (Sec. 3.3.2).

Along with the set of 3D hand joint positions, our 1D convolution-based *ConNet* $f(\cdot)$ also estimates the contact probabilities $\mathbf{b} \in \mathbb{R}^{N \times l}$ on the hand vertices from the hand joint and object pose sequence with a conditioning vector \mathbf{c} that consists of a mass value m and an action label \mathbf{a} .

ConNet $f(\cdot)$ is trained using a binary cross entropy (BCE) loss with the GT hand contact labels l_{con} :

$$\mathcal{L}_{\text{con.}} = \text{BCE}(f(\mathbf{J}^{(0)}, \Phi^{(0)}, \mathbf{c}), l_{\text{con.}}), \quad (8)$$

where $\mathbf{J}^{(0)}$ and $\Phi^{(0)}$ denotes a set of GT 3D hand joints and GT object poses, respectively. At test time, *ConNet* estimates the contact probabilities from the synthesized 3D hand joints and object positions conditioned on \mathbf{c} . The estimated contact probabilities \mathbf{b} are used in the subsequent *fitting optimization* step, to increase the plausibility of the hand and object interactions.

The objective \mathcal{L}_H for the training of *HandDiff* reads:

$$\mathcal{L}_H = \mathcal{L}_{\text{simple}} + \lambda_{\text{geo}}\mathcal{L}_{\text{geo}}, \quad (9)$$

where $\mathcal{L}_{\text{simple}}$ is computed following Eq. (6) and

$$\mathcal{L}_{\text{geo}} = \lambda_{\text{rec.}}\mathcal{L}_{\text{rec.}} + \lambda_{\text{vel.}}\mathcal{L}_{\text{vel.}} + \lambda_{\text{acc.}}\mathcal{L}_{\text{acc.}} + \lambda_{\text{blen.}}\mathcal{L}_{\text{blen.}}. \quad (10)$$

$\mathcal{L}_{\text{rec.}}$, $\mathcal{L}_{\text{vel.}}$ and $\mathcal{L}_{\text{acc.}}$ are loss terms to penalize the positions, velocities, and accelerations of the synthesized hand joints, respectively:

$$\mathcal{L}_{\text{rec.}} = \|\hat{\mathbf{J}}^{(0)} - \mathbf{J}^{(0)}\|_2^2, \quad (11)$$

$$\mathcal{L}_{\text{vel.}} = \|\hat{\mathbf{J}}_{\text{vel.}}^{(0)} - \mathbf{J}_{\text{vel.}}^{(0)}\|_2^2, \quad (12)$$

$$\mathcal{L}_{\text{acc.}} = \|\hat{\mathbf{J}}_{\text{acc.}}^{(0)} - \mathbf{J}_{\text{acc.}}^{(0)}\|_2^2, \quad (13)$$

where $\hat{\mathbf{J}}^{(0)}$ is an approximated set of hand joints from Eq. (7) and $\mathbf{J}^{(0)}$ denotes a set of GT hand joints. $\hat{\mathbf{J}}^{(0)}$ and $\mathbf{J}^{(0)}$ with the subscripts “vel.” and “acc.” represent the velocities and accelerations computed from their positions, respectively.

$\mathcal{L}_{\text{blen.}}$ penalizes incorrect bone lengths of the hand joints using the function $d_{\text{blen.}} : \mathbb{R}^{N \times 3K} \rightarrow \mathbb{R}^{N \times K}$ that computes bone lengths of hands given a sequence 3D hand joints of N frames:

$$\mathcal{L}_{\text{blen.}} = \|d_{\text{blen.}}(\hat{\mathbf{J}}^{(0)}) - d_{\text{blen.}}(\mathbf{J}^{(0)})\|_2^2. \quad (14)$$

At test time, we obtain a set of 3D hand joints \mathbf{J}^* using the denoising process detailed in Eq. (5) given a Gaussian noise $\sim N(0, \mathbf{I})$.

Fitting Optimization Once the 3D hand joint sequence \mathbf{J}^* is synthesized from the trained \mathcal{H} , we solve an optimization problem to fit GHUM hand models to \mathbf{J}^* . We use a threshold of $\mathbf{b} > 0.5$ to select the effective contacts from the per-vertex contact probability obtained in the previous step. Let $\mathbf{b}_{\text{idx}}^n \subset \llbracket 1, L \rrbracket$ be the subset of hand vertex indices with effective contacts on the n -th frame. The objectives are written as follows:

$$\underset{\tau, \phi, \theta}{\text{argmin}} (\lambda_{\text{data}}\mathcal{L}_{\text{data}} + \lambda_{\text{touch}}\mathcal{L}_{\text{touch}} + \lambda_{\text{col.}}\mathcal{L}_{\text{col.}} + \lambda_{\text{prior}}\mathcal{L}_{\text{prior}}). \quad (15)$$

$\mathcal{L}_{\text{data}}$ is a data term to minimize the Euclidean distances between the GHUM (\mathbf{J}) and the synthesized hand joint key points (\mathbf{J}^*):

$$\mathcal{L}_{\text{data}} = \|\mathbf{J} - \mathbf{J}^*\|_2^2. \quad (16)$$

$\mathcal{L}_{\text{touch}}$ is composed of two terms. The first term reduces the distances between the contact hand vertices and their nearest vertices \mathbf{P} on the object to improve the plausibility of the interactions. The second term takes into account the

normals of the object and hands which also enhances the naturalness of the grasp by minimizing the cosine similarity $s(\cdot)$ between the normals of the contact hand vertices \mathbf{n} and the normals of their nearest vertices of the object $\hat{\mathbf{n}}$.

$$\mathcal{L}_{\text{touch}} = \sum_{i=1}^N \sum_{j \in \mathbf{b}_{\text{idx}}^i} \|\mathbf{V}_i^j - \mathbf{P}_i^j\|_2^2 + \sum_{i=1}^N \sum_{i \in \mathbf{b}_{\text{idx}}} (1 - s(\mathbf{n}_i^j, \hat{\mathbf{n}}_i^j)), \quad (17)$$

where the subscript i denotes i -th sequence frame and the superscript j denotes the index of the vertex with the effective contact. $\mathcal{L}_{\text{col.}}$ reduces the collisions between the hand and object by minimizing the penetration distances. Let $\mathcal{P}^n \subset \llbracket 1, U \rrbracket$ be the subset of hand vertex indices with collisions on n -th frame. Then we define

$$\mathcal{L}_{\text{col.}} = \sum_{i=1}^N \sum_{j \in \mathcal{P}^n} \|\mathbf{V}_i^j - \mathbf{P}_i^j\|_2^2. \quad (18)$$

$\mathcal{L}_{\text{prior}}$ is a hand pose prior term that encourages the plausibility of the GHUM hand pose by minimising the pose vector θ of the GHUM parametric model

$$\mathcal{L}_{\text{prior}} = \|\theta\|_2^2. \quad (19)$$

With all these loss terms combined, our final output shows a highly plausible hand and object interaction sequence. The effectiveness of the loss terms is shown in our ablative study (Sec. 4.1). Note that only for the non-spherical objects, which were not present in the training dataset, we apply a Gaussian smoothing on the hand and object vertices along the temporal direction with a sigma value of 3 after the fitting optimization to obtain a smoother motion.

3.3. Object Trajectory Generation

The input object trajectory for *HandDiff* can be provided in two ways, (1) synthesizing 3D trajectory by *TrajDiff* (Sec.3.3.1) or (2) providing a manual trajectory (Sec. 3.3.2). The former allows generating an arbitrary number of hands-object interaction motions conditioned on mass values and action labels, which can contribute to a large-scale dataset generation for machine learning applications. The latter allows for tighter control of the synthesized motions which are still conditioned on an object’s mass value but restricted to the provided trajectory.

3.3.1 Object Trajectory Synthesis

To provide a 3D object trajectory to *HandDiff*, we introduce a diffusion model-based architecture *TrajDiff* that synthesizes an object trajectory given a mass value m and an action label $\mathbf{a} \in \mathbb{R}^6$ encoded as a one-hot vector. We observed that directly synthesizing a set of object rotation values causes jitter artifacts. We hypothesize that this issue comes

from simultaneously synthesizing two aspects of a pose, translation and rotation, each having a different representation. As a remedy, we propose to represent both the translation and rotation as 3D coordinates in a Cartesian coordinate system. Specifically, we first synthesize the *reference vertex positions* \mathbf{P}_{ref} on the object surface defined in the object reference frame, and register them to the predefined *template vertex positions* \mathbf{P}_{temp} to obtain the rotation of the object. We define 6 template vertices as shown in Fig. 3. *TrajDiff* thus synthesizes a set of reference vertex positions $\mathbf{P}_{\text{ref}} \in \mathbb{R}^{N \times q}$ where $q = 18 (= 6 \times 3)$ that are defined in the object center frame along with a set of global translations. We then apply Procrustes alignment between \mathbf{P}_{ref} and \mathbf{P}_{temp} to obtain the object rotations. The objective of *TrajDiff* is defined as follows:

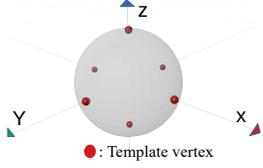


Figure 3. Definition of the template vertices.

$$\mathcal{L}_{\mathcal{T}} = \mathcal{L}_{\text{simple}} + \lambda_{\text{geo.}}(\lambda_{\text{rec.}}\mathcal{L}_{\text{rec.}} + \lambda_{\text{vel.}}\mathcal{L}_{\text{vel.}} + \lambda_{\text{acc.}}\mathcal{L}_{\text{acc.}} + \lambda_{\text{ref.}}\mathcal{L}_{\text{ref.}}). \quad (20)$$

$\mathcal{L}_{\text{rec.}}$, $\mathcal{L}_{\text{vel.}}$ and $\mathcal{L}_{\text{acc.}}$ follow the definitions given in Eqs. (11), (12) and (13), where $\mathbf{J}^{(0)}$ is replaced with GT 3D object poses whose rotation is represented by the reference vertex positions instead of 6D rotation. \mathcal{L}_{ref} is defined as:

$$\mathcal{L}_{\text{ref}} = \|\hat{\mathbf{P}}_{\text{ref}}^{(0)} - \mathbf{P}_{\text{ref}}^{(0)}\|_2^2 + \|d_{\text{rel}}(\hat{\mathbf{P}}_{\text{ref}}^{(0)}) - d_{\text{rel}}(\mathbf{P}_{\text{ref}}^{(0)})\|_2^2. \quad (21)$$

The first term of \mathcal{L}_{ref} penalizes the Euclidean distances between the approximated reference vertex positions $\hat{\mathbf{P}}_{\text{ref}}^{(0)}$ of Eq. (7) and the GT reference vertex positions $\mathbf{P}_{\text{ref}}^{(0)}$. The second term of \mathcal{L}_{ref} penalizes the incorrect Euclidean distances of the approximated reference vertex positions relative to each other. To this end, we use a function $d_{\text{rel}} : \mathbb{R}^{N \times 3q} \rightarrow \mathbb{R}^{N \times q'}$, where $q' = \binom{q}{2}$, which computes the distances between all the input vertices pairs on each frame.

The generated object trajectory responds to the specified masses. Thus, the motion range and the velocity of the object tend to be larger for smaller masses. In contrast, with a heavier object the trajectory shows slower motion and a more regulated motion range.

3.3.2 User-Provided Object Trajectory

Giving the user control over the output in synthesis tasks is crucial for downstream applications such as character animations or avatar generation. Thanks to the design of our architecture that synthesizes 3D hand motions and hand contacts from a mass value and an object trajectory, a manually drawn object trajectory can also be provided to our framework as an input.

However, manually drawing an input 3D trajectory is not straightforward, as it must consider the object dynamics influenced by the mass. For instance, heavy objects will accelerate and/or decelerate much slower than lighter ones. Drawing such trajectories is tedious and often requires professional manual labour. To tackle this issue, we introduce a module that accepts a (user-specified) trajectory with an arbitrary number of points along with the object’s mass, and outputs a *normalized target trajectory* (NTT).

NTT is calculated from the input trajectory based on the intermediate representation that we call *vector of ratios*, see our supplementary for its overview. First, the input (user-specified) trajectory is re-sampled uniformly to $N_{\text{fix}} = 720$ points and passed to *RatioNet*, which for each time step estimates the distance traveled along the trajectory normalized to the range $[0, 1]$ (e.g. the value of 0.3 means that the object traveled 30% of the full trajectory within the given time step). The NTT from this stage is further sent to the *Hand Motion Synthesis* stage to obtain the final hand and object interaction motions. We next explain 1) the initial uniform trajectory re-sampling and 2) the intermediate ratio updates. *Uniform Input Trajectory Re-sampling.* To abstract away the variability of the number of points in the user-provided trajectory of N_{user} points, we first interpolate it into a path Φ_{fix} of length N_{fix} points. Note that N_{user} is not fixed and can vary. We also compute the total path length d_{user} that is used as one of the inputs to the *RatioNet* network (elaborated in the next paragraph):

$$d_{\text{user}} = \sum_{i=1}^{N_{\text{fix}}-1} \|\Phi_{\text{fix}}^i - \Phi_{\text{fix}}^{i+1}\|^2, \quad (22)$$

where Φ_{fix}^i denotes the i -th object position in Φ_{fix} .

Intermediate Ratio Updates Estimation. From the normalized object path Φ_{fix} , a total distance of the path d_{user} , and mass m , we obtain the information of the object location in each time step using a learning-based approach. To this end, we introduce a MLP-based network *RatioNet* $R(\cdot)$ that estimates the location of the object along the path Φ_{fix} encoded as a ratio starting from the beginning, see our supplementary for the schematic visualization. Specifically, *RatioNet* accepts the residual of Φ_{fix} denoted as $\bar{\Phi}_{\text{fix}}$, a mass scalar value and d_{user} and outputs a vector $\mathbf{r} \in \mathbb{R}^N$ that contains the update of the ratios on the path for each time step:

$$\mathbf{r} = R(\bar{\Phi}_{\text{fix}}, m, d_{\text{user}}). \quad (23)$$

Next, we obtain the cumulative ratios \mathbf{r}_{cuml} from \mathbf{r} starting from the time step 0 to the end of the frame sequence. Finally, the NTT $\Phi_{\text{NTT}} = [\Phi_{\text{NTT}}^0, \dots, \Phi_{\text{NTT}}^N]$ at time step t is obtained as:

$$\Phi_{\text{NTT}}^t = \Phi_{\text{fix}}^{id}, \quad \text{with } id = \text{round}(r_{\text{cuml}}^t \cdot N_{\text{fix}}), \quad (24)$$

where id and “.” denote the index of Φ_{fix} , and multiplication, respectively. *RatioNet* is trained with the following loss function \mathcal{L}_{ratio} :

$$\mathcal{L}_{ratio} = \|\mathbf{r} - \hat{\mathbf{r}}\|_2^2 + \|\mathbf{r}_{vel} - \hat{\mathbf{r}}_{vel}\|_2^2 + \|\mathbf{r}_{acc} - \hat{\mathbf{r}}_{acc}\|_2^2 + \mathcal{L}_{one}, \quad (25)$$

$$\mathcal{L}_{one} = \left\| \left(\sum_{i=1}^N \mathbf{r}^i \right) - \mathbf{1} \right\|_2^2, \quad (26)$$

where $\hat{\mathbf{r}}$ denotes the GT ratio updates. Note that all terms in Eq. (25) have the same weights. The subscripts “vel.” and “acc.” represent the velocity and accelerations of \mathbf{r} and $\hat{\mathbf{r}}$, respectively. \mathcal{L}_{one} encourages *RatioNet* to estimate the sum of the ratio updates to be 1.0.

4. Experiments

To the best of our knowledge, there exists no other work that addresses the hand object manipulation synthesis conditioned on mass. Therefore, we compare our method mainly with two baseline methods which, similarly to our method, employ an encoder-decoder architecture, but which are based on the popular methods VAE [16] and VAEGAN [39]. Specifically, the VAE baseline uses the same diffusion model architecture as our method, but we add a reparameterization layer [16] and remove the skip connections between the encoder and the decoder. The VAEGAN baseline shares the same architecture of the generator, while the discriminator network consists of three 1D convolution layers and two fully connected layers at the output of the network, and we use ELU activation in the discriminator [5]. The generator and discriminator networks are conditioned by the same conditioning vector. In all the following experiments we will refer to our proposed method as *Ours* and to the baselines as *VAE* and *VAEGAN*. We also compare with ManipNet [41] qualitatively, while the quantitative comparison is omitted due to the following limitations of ManipNet. (1) It requires a sequence of 6D hand and object poses as inputs, whereas our approach only needs conditioning of mass value and an optional action label, (2) certain evaluation metrics (e.g., diversity, multimodality) cannot be fairly computed on ManipNet due to its deterministic nature, and (3) ManipNet lacks control over the object weight as it does not support mass conditioning. Therefore, we compare qualitatively with ManipNet by inputting the ground truth 6D object and hand poses to the method. Please refer to our supplementary material for additional quantitative experiments (additional ablations, qualitative results, and a user study).

4.1. Quantitative Results

In this section, we evaluate the motion quality of *MACS* from various perspectives. We report a diversity and multimodality measurement as suggested by Guo *et al.* [12] in

Table 1. We also evaluate the physical plausibility by measuring the following metrics:

Non-collision ratio (m_{col}) measures the ratio of frames with no hand-object collisions. A higher value indicates fewer collisions between the hand and the object.

Collision distance (m_{dist}) measures the distance of hand object penetration averaged over all the samples. A lower value indicates low magnitude of the collisions.

Non-touching ratio (m_{touch}) measures the ratio of samples over all the samples where there is no contact between the hand and object. A lower value indicates fewer *floating object* artifacts (*i.e.*, spurious absence of contacts).

Note that to report m_{touch} , we discard throwing motion action labels, as the assumption is that there should be constant contact between the hands and the object. The hand vertices whose nearest distances to the object are lower than a threshold value of $5mm$ are considered contact vertices. Similarly, to compute m_{col} and m_{dist} , the interpenetrations over $5mm$ are considered collisions. To compute the metrics, we generate 500 samples across 6 different action labels.

Diversity and Multimodality Diversity measures the motion variance over all the frames within each action class, whereas multimodality measures the motion variance across the action classes. High diversity and multimodality indicate that the generated samples contain diversified motions. Please refer to Guo *et al.* [12] for more details. We report the diversity and multimodality metrics for the generated hand motions and the object trajectories in Table 1. It is clear that in both cases *Ours* generates much more diversified motions when compared to the baselines, which we attribute to our diffusion model-based architecture. Notably, the generated trajectory samples contain more diversified motions compared with the metrics computed on the GT data.

Physical plausibility We report the physical plausibility measurements in Table 2. *Ours* shows the highest performance across all three metrics m_{col} , m_{dist} and m_{touch} . *VAE* yields m_{col} and m_{dist} comparable to *Ours*, however, its m_{touch} is substantially higher with 42% error increase compared to *Ours*. *VAEGAN* shows m_{touch} similar to *Ours* but it underperforms in terms of the collision-related metrics.

Ablation study Here, we motivate the use of the important loss terms of our fitting optimization and training loss functions. In Table 2, we show the results of the fitting optimization without \mathcal{L}_{touch} and without \mathcal{L}_{col} . When omitting the contact term \mathcal{L}_{touch} , the generated hands are not in contact with the object in most of the frames. This results in substantially higher metric m_{touch} and manifests through undesirable *floating object* artifacts. Omitting the collision term \mathcal{L}_{col} leads to frequent interpenetrations, lower m_{col} and higher m_{dist} . Therefore, it is essential to employ both the loss terms to generate sequences with higher physical

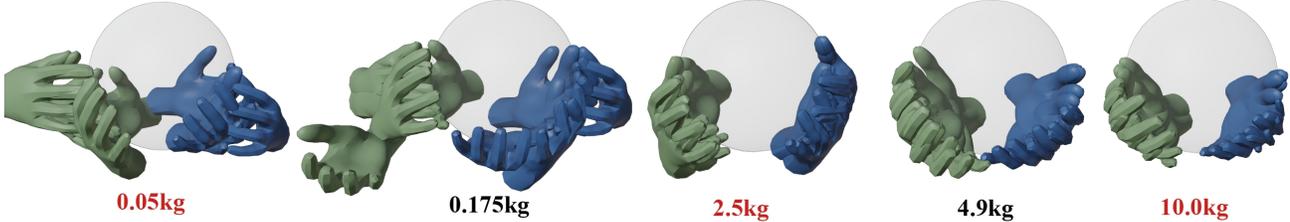


Figure 4. Grasp synthesis with different object masses. Our method can generate sequences influenced by masses close (in black) and far (in red) from the training dataset. Note that in the case of small masses, hands can support the object with fingertips and release the object for some time; the hands are generally more mobile. The situation is different for moderate and large masses: A larger area supporting the object is necessary, and the hands are less mobile.

	Hand synthesis		Trajectory synthesis	
	Diversity \uparrow	Multimodality \uparrow	Diversity \uparrow	Multimodality \uparrow
GT	9.984 \pm 0.36	7.255 \pm 0.32	10.041 \pm 0.28	7.374 \pm 0.29
Ours	9.606\pm0.33	7.07\pm0.30	11.01\pm0.37	8.05\pm0.33
VAE	8.350 \pm 0.42	6.0465 \pm 0.34	9.584 \pm 0.47	7.696 \pm 0.43
VAEGAN	7.821 \pm 0.27	5.887 \pm 0.26	8.428 \pm 0.29	6.285 \pm 0.30

Table 1. Diversity and multimodality for the hand and trajectory synthesis compared to the ground truth.

	m_{col} [%] \uparrow	m_{dist} [mm] \downarrow	m_{touch} [%] \downarrow
Ours	97.84	0.041	1.97
Ours w/o $\mathcal{L}_{\text{touch}}$	100.0	0.0	63.3
Ours w/o \mathcal{L}_{col}	38.41	0.296	1.88
VAE	97.2	0.055	2.80
VAE-GAN	96.03	0.058	2.03

Table 2. Physical plausibility measurement of our full model and its trimmed versions vs VAE and VAE-GAN.

plausibility. For more ablations for the loss terms \mathcal{L}_{vel} and \mathcal{L}_{acc} for the network training, also for ablation on RatioNet, please refer to our supplementary material.

4.2. Qualitative Results

Hand-Object Interaction Synthesis In our supplementary video, we show the synthesized hand and object interaction sequence conditioned by the action labels and mass of the object. The synthesized motions show realistic and dynamic interactions between the hands and the object. Furthermore, thanks to our cascaded diffusion models, the generated motions show high diversity. The results thus visually clearly complement the quantitative findings listed in Table 1. Furthermore, our method shows a more robust and plausible synthesis that faithfully responds to the conditioning mass value compared to ManipNet [41].

Grasp Synthesis We show 5 samples of grasps for different conditioning mass values in Fig. 4. To generate this visualization, we trained *HandDiff* without providing the action labels. In order to synthesize the grasps, we provide an object trajectory with position and rotations set to 0. Our method shows diverse grasps faithfully reflecting the conditional mass values. Most notably, the synthesized hands tend to support the heavy object at its bottom using the whole palm, whereas the light object tends to be supported using the fingertips only. Furthermore, the synthesized grasps show reasonable results even with unseen interpolated (2.5kg) and extrapolated (0.05kg and 10.0kg) mass values (highlighted in red).

5. Conclusion

This paper introduces the first approach to synthesize realistic 3D object manipulations with two hands faithfully responding to conditional mass. Our diffusion-model-based *MACS* approach produces plausible and diverse object manipulations, as verified quantitatively and qualitatively.

Since this topic has so far been completely neglected in the literature, the focus of this paper is to demonstrate the impact of mass onto manipulation and hence we opted to use a single shape with uniform static mass distribution. As such there are several limitations that open up to exciting future work; for example the effect of shape diversity, non-uniform mass distribution (i.e. one side of the object is heavier than the other), or dynamic mass distribution (e.g., a bottle of water). Furthermore, we would like to highlight that other physical factors, such as friction or individual muscle strength, also impact object manipulation and could be addressed in future works. Lastly, while this work focused on synthesis with applications for ML data generation, entertainment and mixed reality experiences, we believe that weight analysis is another interesting avenue to explore, i.e. predicting the weight based on observed manipulation. This could be valuable in supervision scenarios to identify if an object changed its weight over time.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian

- Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. 12
- [2] G. Bradski. The OpenCV Library. *Dr. Dobbs's Journal of Software Tools*, 2000. 12
- [3] Samarth Brahmabhatt, Cusuh Ham, Charles C Kemp, and James Hays. Contactdb: Analyzing and predicting grasp contact via thermal imaging. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 1
- [4] Sammy Christen, Muhammed Kocabas, Emre Aksan, Jemin Hwangbo, Jie Song, and Otmar Hilliges. D-grasp: Physically plausible dynamic grasp synthesis for hand-object interactions. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2
- [5] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015. 7
- [6] Enric Corona, Albert Pumarola, Guillem Alenya, Francesc Moreno-Noguer, and Grégory Rogez. Ganhand: Predicting human grasp affordances in multi-object scenes. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [7] Rishabh Dabral, Muhammad Hamza Mughal, Vladislav Golyanik, and Christian Theobalt. Mofusion: A framework for denoising-diffusion-based motion synthesis. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 4, 12
- [8] Renaud Detry, Dirk Kraft, Anders Glent Buch, Norbert Krüger, and Justus Piater. Refining grasp affordance models by experience. In *International Conference on Robotics and Automation (ICRA)*, 2010. 1
- [9] Anindita Ghosh, Rishabh Dabral, Vladislav Golyanik, Christian Theobalt, and Philipp Slusallek. Imos: Intent-driven full-body motion synthesis for human-object interactions. In *Eurographics*, 2023. 1, 2
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014. 3
- [11] Patrick Grady, Chengcheng Tang, Christopher D Twigg, Minh Vo, Samarth Brahmabhatt, and Charles C Kemp. Contactopt: Optimizing contact to improve grasps. In *Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2
- [12] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *ACM International Conference on Multimedia*, 2020. 7
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2, 3, 4, 12
- [14] Haoyu Hu, Xinyu Yi, Hao Zhang, Jun-Hai Yong, and Feng Xu. Physical interaction: Reconstructing hand-object interactions with physics. In *SIGGRAPH Asia 2022 Conference Papers*, 2022. 1
- [15] Korrawe Karunratanakul, Jinlong Yang, Yan Zhang, Michael J Black, Krikamol Muandet, and Siyu Tang. Grasping field: Learning implicit representations for human grasps. In *International Conference on 3D Vision (3DV)*, 2020. 1, 2
- [16] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations (ICLR)*, 2014. 7
- [17] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations (ICLR)*, 2021. 2
- [18] Robert Krug, Dimitar Dimitrov, Krzysztof Charusta, and Boyko Iliev. On the efficient computation of independent contact regions for force closure grasps. In *International Conference on Intelligent Robots and Systems (ICIRS)*, 2010. 1, 2
- [19] Ying Li, Jiaxin L Fu, and Nancy S Pollard. Data-driven grasp synthesis using shape matching and task-based pruning. *Transactions on visualization and computer graphics (TVCG)*, 13(4):732–747, 2007. 1, 2
- [20] Shaowei Liu, Hanwen Jiang, Jiarui Xu, Sifei Liu, and Xiaolong Wang. Semi-supervised 3d hand-object poses estimation with interactions in time. In *Computer Vision and Pattern Recognition (CVPR)*, 2021. 1
- [21] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chu-Ling Chang, Ming Yong, Juhyun Lee, et al. Mediapipe: A framework for perceiving and processing reality. In *Workshop on Computer Vision for AR/VR at Computer Vision and Pattern Recognition (CVPRW)*, 2019. 11
- [22] Igor Mordatch, Zoran Popović, and Emanuel Todorov. Contact-invariant optimization for hand manipulation. In *Proceedings of the ACM SIGGRAPH/Eurographics symposium on computer animation (SCA)*, pages 137–144, 2012. 1
- [23] Franziska Mueller, Micah Davis, Florian Bernard, Oleksandr Sotnychenko, Mickeal Verschoor, Miguel A Otaduy, Dan Casas, and Christian Theobalt. Real-time pose and shape reconstruction of two interacting hands with a single depth camera. *ACM Transactions on Graphics (TOG)*, 38(4), 2019. 1
- [24] Nancy S Pollard and Victor Brian Zordan. Physically based grasping control from example. In *ACM SIGGRAPH/Eurographics symposium on Computer animation (SCA)*, 2005. 1, 2
- [25] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *International Conference on Learning Representations (ICLR)*, 2023. 2
- [26] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 4

- [27] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in neural information processing systems (NeurIPS)*, 2022. 2
- [28] Matthias Schroder and Helge Ritter. Hand-object interaction detection with fully convolutional networks. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [29] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning (ICML)*, 2015. 2, 3
- [30] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems (NeurIPS)*, 2015. 3
- [31] Omid Taheri, Nima Ghorbani, Michael J. Black, and Dimitrios Tzionas. GRAB: A dataset of whole-body human grasping of objects. In *European Conference on Computer Vision (ECCV)*, 2020. 1
- [32] Omid Taheri, Vasileios Choutas, Michael J Black, and Dimitrios Tzionas. Goal: Generating 4d whole-body motion for hand-object grasping. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [33] Bugra Tekin, Federica Bogo, and Marc Pollefeys. H+ o: Unified egocentric recognition of 3d hand-object poses and interactions. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 1
- [34] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. In *International Conference on Learning Representations (ICLR)*, 2023. 2, 3
- [35] Anand Thobbi and Weihua Sheng. Imitation learning of hand gestures and its evaluation for humanoid robots. In *International Conference on Information and Automation (ICIA)*, 2010. 2
- [36] Jiayi Wang, Franziska Mueller, Florian Bernard, Suzanne Sorli, Oleksandr Sotnychenko, Neng Qian, Miguel A Otaduy, Dan Casas, and Christian Theobalt. Rgb2hands: real-time tracking of 3d hand interactions from monocular rgb video. *ACM Transactions on Graphics (TOG)*, 39(6), 2020. 1
- [37] Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Ghum & ghuml: Generative 3d human shape and articulated pose models. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. 3, 11
- [38] Yuting Ye and C Karen Liu. Synthesis of detailed hand manipulations using contact sampling. *ACM Transactions on Graphics (ToG)*, 31(4), 2012. 1, 2
- [39] Xianwen Yu, Xiaoning Zhang, Yang Cao, and Min Xia. Vae-gan: A collaborative filtering framework based on adversarial variational autoencoders. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2019. 7
- [40] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. In *International Conference on Computer Vision (ICCV)*, 2023. 2
- [41] He Zhang, Yuting Ye, Takaaki Shiratori, and Taku Komura. Manipnet: Neural manipulation synthesis with a hand-object spatial representation. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021. 1, 2, 7, 8
- [42] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. 2
- [43] Juntian Zheng, Qingyuan Zheng, Lixing Fang, Yun Liu, and Li Yi. Cams: Canonicalized manipulation spaces for category-level functional hand-object manipulation synthesis. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [44] Keyang Zhou, Bharat Bhatnagar, Jan Eric Lenssen, and Gerard Pons-Moll. Toch: Spatio-temporal object correspondence to hand for motion refinement. In *European Conference on Computer Vision (ECCV)*, 2022. 1, 2
- [45] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 3

Appendices

This supplementary document provides the details of our dataset acquisition (Sec. A), network architectures (Sec. B), and implementations (Sec. C). We also provide further ablations (1) for the loss terms $\mathcal{L}_{vel.}$ and $\mathcal{L}_{acc.}$ for the network training, (2) for the mass conditioning and (3) for ablation on RatioNet and (4) user-study on the synthesized motions. We also show additional qualitative results for (5) the objects unseen during the training, (6) visualizations of the synthesized contacts and (7) the synthesized motions given a user-provided trajectory (Sec. D).

A. Dataset

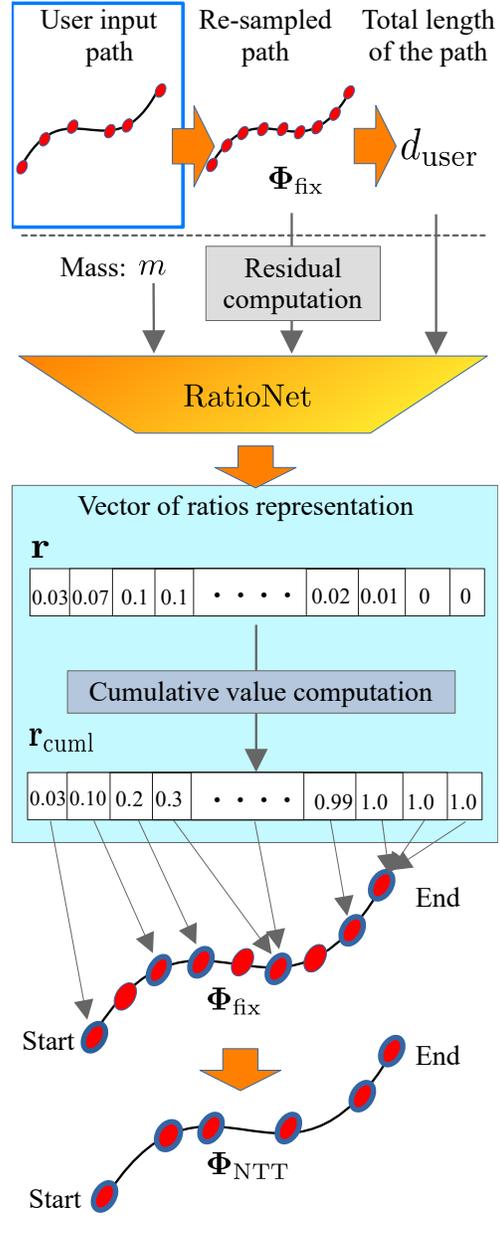
Since there exists no 3D hand and object interaction motion dataset with corresponding object mass values of the objects, we reconstruct such motions using 8 synchronized Z-CAM E2 cameras of 4K resolution and 50 fps. As target objects, we use five plastic spheres of the same radius 0.1[m]. We fill them with different materials of different densities to prepare the objects of the same volume and different weights *i.e.* 0.175, 2.0, 3.6, 3.9, 4.9 kg. Each sphere is filled entirely so that its center of mass does not shift as the object is moved around. Five different subjects are asked to perform five different actions manipulating the object: (1) vertical throw and catch, (2) passing from one hand to another, (3) lifting up and down, (4) moving the object horizontally, and (5) drawing a circle. The subjects perform each action using both their hands while standing in front of the cameras and wearing colored wristbands (green for the right wrist and yellow for the left wrist), which are later used to classify handedness. The recordings from the multi-view setup were further used to reconstruct the 3D hand and object motions, totaling 110k frames. The details of the capture and reconstruction processes are described in the following text.



Figure 5. Image of our markeded sphere and recording example.

Five different subjects are asked to perform five different actions manipulating the object: (1) vertical throw and catch, (2) passing from one hand to another, (3) lifting up and down, (4) moving the object horizontally, and (5) drawing a circle. The subjects perform each action using both their hands while standing in front of the cameras and wearing colored wristbands (green for the right wrist and yellow for the left wrist), which are later used to classify handedness. The recordings from the multi-view setup were further used to reconstruct the 3D hand and object motions, totaling 110k frames. The details of the capture and reconstruction processes are described in the following text.

Hand Motion Reconstruction To reconstruct 3D hand motions, we first obtain 2D hand key points from all the camera views using MediaPipe [21]. We then fit GHUM hand models [37] for both hands on each frame by solving 2D keypoint reprojection-based optimization with the known camera intrinsics and extrinsic combining with a collision loss term (Eq.(18)), a pose prior loss (Eq.(19)) in our main paper and a shape regularizer term that minimizes the norm of the shape parameter β of the GHUM hand parametric model.



Notations

- Φ_{fix} : Uniformly re-sampled user-specified trajectory
- \mathbf{r} : Update of the ratios on the path for each time step
- \mathbf{r}_{cuml} : Cumulative values of \mathbf{r}
- Φ_{NTT} : Normalized target trajectory

Figure 6. Schematic visualization of the user input trajectory processing stage.

	Ours	Ours w/o \mathcal{L}_{vel}	Ours w/o \mathcal{L}_{acc}
acc. dist. ↓	7.35	26.4	11.2

Table 3. Wasserstein distances between the acceleration distributions (“acc. dist”) of the generated motions and ground-truth motions. Combining both \mathcal{L}_{vel} and \mathcal{L}_{acc} shows the highest plausibility in terms of the accelerations.

	0.175 [kg]	2.0 [kg]	3.6 [kg]	3.9 [kg]	4.9 [kg]
ours	0.006	0.010	0.012	0.011	0.011
ours w/o cond.	0.089	0.070	0.081	0.061	0.074

Table 4. Wasserstein distances between the acceleration distributions (“acc. dist”) of the generated and ground-truth motions.

Object Trajectory Reconstruction We place around 50 ArUco markers of the size 1.67×1.67 cm on each sphere for the tracking optimization (see Fig. 5 for the example of our tracking object). The marker positions in the image space are tracked using the OpenCV [2] library. The 3D object positions on each frame are obtained by solving the multi-view 2D reprojection-based optimization.

B. Network Architecture

We employ the Unet-based diffusion model networks from Ho et al. [13] for our *TrajDiff* and *HandDiff*. *HandDiff* uses four sets of 2D convolutional residual blocks for the encoder and decoder architecture. *TrajDiff* is composed of two sets of residual blocks of 1D convolution layers instead of 2D convolutions. The number of kernels at its output 1D convolutional layer is set to 21 which corresponds to the dimensionality of the object pose. *ConNet* consists of three-1D convolutional layers with ELU and a sigmoid activation for its hidden layers and output layer, respectively. Similarly, *RatioNet* is composed of three-layer MLP with ELU and a sigmoid activation functions in the hidden and output layers, respectively. Starting from the input layer, the output layer dimensions are 1024, 512 and 180. See Fig. 6 for the overview of the user input trajectory processing stage (Sec. 3.3.2 in the main paper) that utilizes *RatioNet*.

C. Training and Implementation Details

All the networks are implemented in TensorFlow [1] and trained with 1 GPU Nvidia Tesla V100 until convergence. The training of *HandDiff*, *TrajDiff*, *ConNet* and *RatioNet* takes 5 hours, 3 hours, 2 hours and 2 hours, respectively. We set the loss term weights of Eq. (10) and (20) to $\lambda_{rec.} = 1.0$, $\lambda_{vel.} = 5.0$ and $\lambda_{acc.} = 5.0$. $\lambda_{blen.}$ of Eq. (10) and λ_{ref} of Eq. (20) are set to 10.0 and 1.0, respectively. For the fitting optimization defined in Eq. (15), we set $\lambda_{data} = 1.0$, $\lambda_{touch} = 0.7$, $\lambda_{col.} = 0.8$ and $\lambda_{prior} = 0.001$. As in Dabral

	Ours	Interpolation
acc. dist. ↓	0.379	0.447

Table 5. Wasserstein distances between the acceleration distributions (“acc. dist”) of ground-truth trajectory and the generated from *RatioNet* (Ours). We also show the same metric computed on the interpolated subdivided trajectory with an equal length.

	GT	Ours	VAE	VAEGAN
reality score ↑	7.10±2.09	6.01±2.08	5.10±2.24	4.54±2.39

Table 6. Results of the user study (perceptual motion quality).

et al. [7], $\lambda_{geo.}$ of Eq. (10) and (20) are set such that larger penalties are applied with smaller diffusion steps t :

$$\lambda_{geo.} = \frac{10}{\exp \frac{10t}{T}}, \quad (27)$$

where T is the maximum diffusion step. We empirically set the maximum diffusion step T for *HandDiff* and *TrajDiff* to 150 and 300, respectively.

D. Further Evaluations

In this section, we show further ablative studies to evaluate the significance of the components in our method.

Temporal loss terms $\mathcal{L}_{vel.}$ and $\mathcal{L}_{acc.}$: to report the ablative study of the loss terms $\mathcal{L}_{vel.}$ and $\mathcal{L}_{acc.}$ for the network training, we compute the Wasserstein distance between the accelerations of the sampled data and the GT data denoted as “acc. dist.” in Table 3. Combining the two loss terms $\mathcal{L}_{vel.}$ and $\mathcal{L}_{acc.}$, our method shows the shortest distance from the GT acceleration distributions.

Plausibility of the conditioning mass value effect: can be evaluated by measuring the similarity between the GT object accelerations and the sampled ones. In Table 4, we show the “acc. dist.” between the accelerations of the ground truth object motions and the sampled motions *with and without* mass conditioning. With the conditioning mass value, our network synthesizes the motions with more physically plausible accelerations on each mass value compared with the network without mass conditioning.

Effect of RatioNet on the user-provided trajectories: The goal of RatioNet is to provide plausible dynamics on the user-provided trajectories given conditioning mass values *e.g.* higher object motion speed appears with lighter mass and the object is moved slower with heavier mass value. For the ablative study of RatioNet, we report the “acc. dist.” with and without RatioNet comparing with the acceleration distributions of our GT trajectories. For the component without RatioNet, we simply apply uniform sampling on the provided trajectories, denoted as “Interpolation” in Table 5.

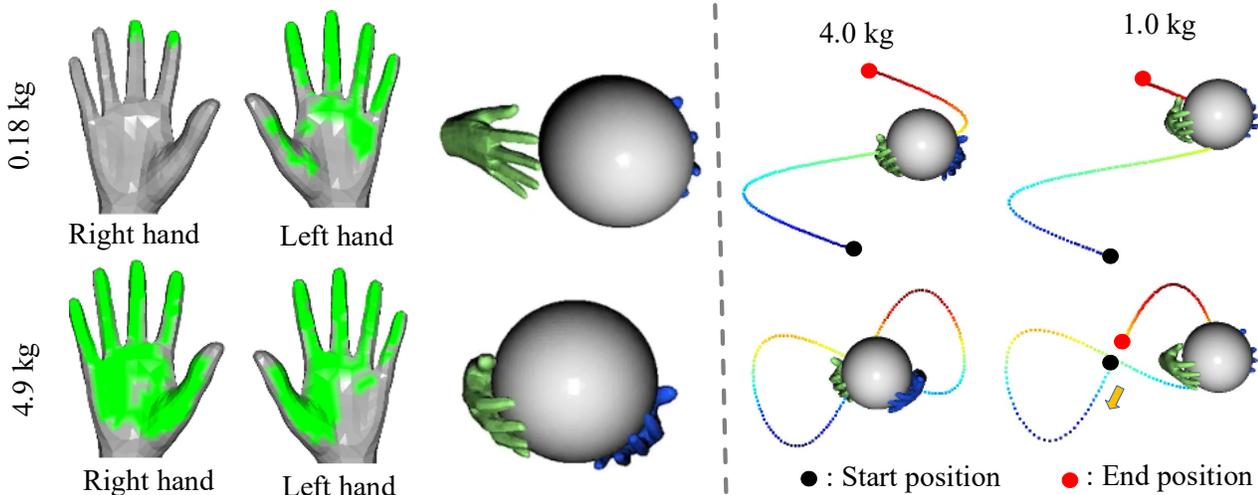


Figure 7. **(left)** Example visualizations of the contacts synthesized by *ConNet*, given conditioning mass values of 0.18 kg (top) and 4.9 kg (bottom). With heavier mass, the contact region spans the entire palm region whereas contacts concentrate around the fingertips for a light object. **(right)** Example visualizations of 3D object manipulation given user input trajectories of S curve (top) and infinity curve (bottom). Thanks to the *RatioNet*, the object manipulation speed matches our intuition *i.e.* slower manipulation speed with heavier objects, and vice versa. See our supplementary video for the sequential visualizations.

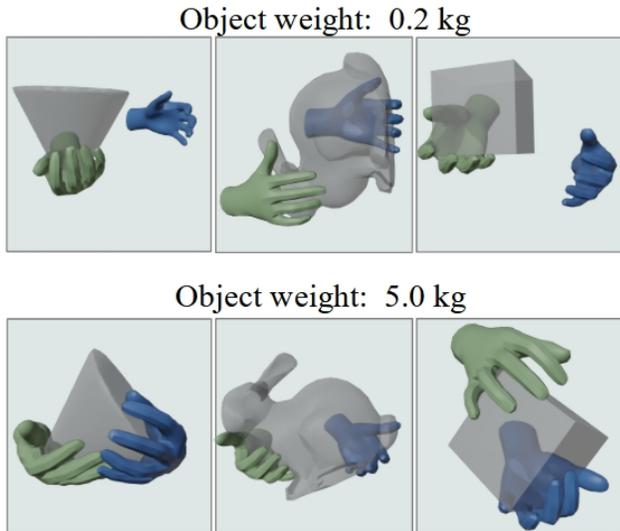


Figure 8. Example visualizations of 3D manipulations of the objects unseen during the training, given conditioning mass value of 0.2kg (top) and 5.0kg (bottom). *MACS* adapts to unseen shapes thanks to its mass-conditioned synthesized hand contacts.

Thanks to our *RatioNet*, the object acceleration shows much more plausible values than without the network, faithfully responding to the conditioning mass values. The qualitative results of *RatioNet* can be seen in our supplementary video.

User Study The realism of 3D motions can be perceived differently depending on individuals. To quantitatively measure the plausibility of the synthesized motions, we perform an online user study. We prepared 26 questions with videos and gathered 42 participants in total. The videos for the study were randomly selected from the sampled results of VAE and VAEGAN baselines, *MACS* and the GT motions. In the first section, the subjects were asked to select the naturalness of the motions on a scale of 1 to 10 *reality score* (1 for completely unnatural and 10 for very natural). Table 6 shows the mean scores. *MACS* clearly outperforms other benchmarks in this perceptual user study, thanks to our diffusion-based networks that synthesize 3D manipulations with high-frequency details. In the additional section, we further evaluated our method regarding how faithfully the synthesized motions are affected by the conditional mass value. We show two videos of motions at once where the network is conditioned by mass values of 1.0 and 5.0, respectively. The participants were instructed to determine which sequence appeared to depict the manipulation of a heavier object. On average, the participants selected the correct answer with 92.8% accuracy, which suggests that *MACS* plausibly reflects the conditioning mass value in its motion.

Qualitative Results In Fig. 7 - (left), we provide visual examples of synthesized contacts with different mass values (0.18kg and 4.9kg). The synthesized contacts are distributed across the palm region when a heavier mass is given, whereas they concentrate around the fingertips with

a lighter mass, which follows our intuition. Additionally, Fig. 7 - (right) displays example synthesis results with user-provided input trajectories (S-curve and infinity curve). Thanks to the RatioNet, the object speed reflects the conditioning mass value, *i.e.* faster speed for lighter mass and vice versa. See our supplementary video for its sequential visualizations.

Unseen Objects In Fig. 8, we show the synthesized motions for objects that were not seen during the training, specifically a cone, the Stanford bunny and a cube. Thanks to the synthesized hand contact labels conditioned by a mass value, *MACS* shows modest adaptations to different shapes while still correctly reflecting the provided mass values.