

FORCE: Dataset and Method for Intuitive Physics Guided Human-object Interaction

Xiaohan Zhang^{1,2}, Bharat Lal Bhatnagar³, Sebastian Starke³,
Ilya Petrov^{1,2}, Vladimir Guzov^{1,2}, Helisa Dharmo⁴,
Eduardo Pérez-Pellitero⁴, and Gerard Pons-Moll^{1,2}

¹ University of Tübingen and Tübingen AI Center

² Max Planck Institute for Informatics, Germany

³ Meta Reality Labs Research

⁴ Huawei Noah's Ark Lab

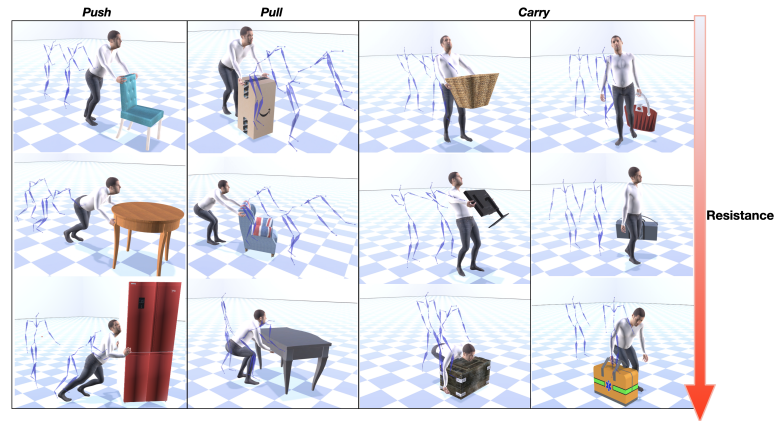


Fig. 1: FORCE, a method and dataset for synthesizing diverse, nuanced human-object interactions by modeling the physical attributes of the interaction.

Abstract. Interactions between human and objects are influenced not only by the object’s pose and shape, but also by physical attributes such as object mass and surface friction. They introduce important motion nuances that are essential for diversity and realism. Despite advancements in recent kinematics-based methods, this aspect has been overlooked. Generating nuanced human motion presents two challenges. First, it is non-trivial to learn from multi-modal human and object information derived from both the physical and non-physical attributes. Second, there exists no dataset capturing nuanced human interactions with objects of varying physical properties, hampering model development. This work addresses the gap by introducing the FORCE model, a kinematic approach for synthesizing diverse, nuanced human-object interactions by modeling physical attributes. Our key insight is that human motion is dictated by the interrelation between the force exerted by the human and the perceived resistance. Guided by a novel intuitive physics encoding, the model captures the interplay between human

force and resistance. Experiments also demonstrate incorporating human force facilitates learning multi-class motion. Accompanying our model, we contribute the FORCE dataset. It features diverse, different-styled motion through interactions with varying resistances. Our code, dataset, and models will be released at <https://virtualhumans.mpi-inf.mpg.de/force/>.

1 Introduction

Synthesizing realistic motion of human-object interaction presents a significant challenge. The complexity lies in the intricate interplay between humans and objects. While previous work has primarily addressed fundamental aspects of interactions, such as object shape and location, they overlook crucial physical attributes like weight, friction, and human force [67,25,92,70]. Consequently, existing models struggle with distinguishing between carrying an empty suitcase and a full one, or determining if the interaction is feasible. These nuances are essential for diversity and realism. This work aims to fill this gap by leveraging the physical attributes to synthesize nuanced human motion in these scenarios.

Physics-based methods coupled with reinforcement learning [57,56,19,84] demonstrate remarkable generalization for different external forces. However, they face challenges including high complexity, because distinctive policies are usually trained with tailored reward functions for navigation and different interaction tasks. This means a hybrid approach is generally adopted. Additionally, these methods exhibit limitations in providing fine-grained control, for example, the switching between left-handed carrying to using both hands.

Kinematics-based methods for synthesizing human motion are more scalable. It is imperative for real-world applications such as AR/VR and gaming, where a *single* model is adopted for synthesizing long-term, complex interactions. However, earlier kinematic methods either disregard the surrounding environment [46,55,45,42] or focus solely on static objects [25,70,79,92,49,91]. Closest to our goal are methods [67,25,30,40] that model object shapes, but overlook physical attributes of interaction. In reality, humans adapt their motion based on perceived resistance and applied force during an object interaction [47,20]. Take the *pushing* action shown in Figure 1 (left-most column). When pushing around a heavier object, the human applies a greater force, naturally adjusting posture, shifting the center of mass [73,12], and leaning forward to push against the friction. If the resistance exceeds the applied force, the object would not move and the human would have to give up the interaction because it is not feasible. Synthesizing such nuanced motion calls for a method that generalizes well to these physical attributes of the interaction.

Developing a kinematic method presents multiple challenges. Firstly, it is non-trivial to reason about multi-modal human-object information, including different actions, object poses, and crucial physical attributes. The increased complexity makes it challenging to disambiguate similar human poses, leading to motion that lacks diversity and nuanced details. Secondly, when determining the feasibility of the interaction, resistance is not the sole factor. It also depends

on how the human interacts. For instance, a human can more effectively carry a heavier object with two hands than with one. It is shown that naively conditioning on the resistance generates sub-optimal results (see Section 5). Additionally, there is no dataset capturing diverse daily interactions under varying physical conditions. The lack of such data hinders model development and evaluation. Challenges present even when collecting such data, when addressing issues such as object occlusion.

To address these challenges, we introduce FORCE, the first kinematic method for synthesizing human-object interactions that focuses on the nuanced details in human motion. Our method is founded on a pivotal insight: Human motion is dictated by the interrelation between the force exerted by the human and the perceived resistance. Guided by a novel intuitive physics encoding derived from these crucial attributes, our model is able to synthesize a diverse spectrum of interactions. For example, under the category of “carrying”, our model can plausibly generate motions such as carrying an object, carrying an object followed by a need to drop it, or attempting to carry but encountering failure. Additionally, we enable interactive control at run-time, the style of motion can be manipulated not only by varying the resistance of the object, but also by the desired action and contact mode (left, right, two-handed).

Complementing our FORCE model, we present a dataset featuring diverse motion nuances through interactions involving 3-6 levels of resistance. We adopt a customized hybrid tracker comprising of 4 Kinect RGB-D cameras [1] paired with 17 Inertial Measurement Units (IMUs) [54]. Our novel dataset contains 450 motion sequences (192k frames) of pervasive interactions of carrying, pushing and pulling objects. For each frame, we provide high-quality human and object poses. This dataset can serve as a benchmark for various human-object interaction tasks. The overall contributions of our work are:

1. We introduce FORCE, the first kinematic method to synthesize human-object interaction by modeling physical attributes such as resistance and the applied human force.
2. To enable the synthesis of diverse, nuanced human motion, we propose a novel intuitive physics encoding.
3. We present a dataset that accurately captures the daily interactions of pushing, pulling and carrying objects. It features diverse, different-styled interaction motions with varying resistances.
4. We will release our code, dataset, and models to stimulate further research.

2 Related Work

Interactions with static objects. Synthesizing human-object interaction is a long-standing challenge in computer vision. Its evolution spanning from the early stages of non-contextual human motion synthesis [29,2,58,65,72,64,17] to more recent work aimed at predicting static affordances within 3D scenes [43,61]. This research domain primarily focuses on predicting static human entities in a given scene [93,90,26,94,86,27], while adhering to scene constraints.

Recently, despite the substantial advancements in human-object interaction learning, most existing research focuses on human interactions with static objects. Several works [75,74] predict human motion in pre-scanned scenes [26]. They train separate modules for predicting root trajectories initially and subsequently generating full-body poses, conditioned on both the scene and the intended path. Nevertheless, the scale and motion quality of these datasets often limit the realism of generated human motion [76,33,38,49]. In efforts to enhance the quality of motion generation, recent research has concentrated on static objects interactions, such as sitting and lying on chairs or sofas [25,91,92,60,52,95], and more recently reaching motion [37].

Another avenue of exploration centers around full-body grasping synthesis [70,79,4]. More development effort has expanded to simultaneously synthesize human and object motions after grasping an object [22,41]. Specialized efforts are dedicated to synthesizing dexterous hand-object interactions using object motion [71,89,97,59]. Additionally, there exists a parallel work focusing on reinforcement learning (RL) [16,9]. Despite being relevant, these methods consider humans interacting with static objects only. We address more challenging scenarios of human interacting with moving objects.

Interactions with moving objects. Investigation into synthesizing human motion for object manipulation covers both physics-based and kinematic-based approaches. In the domain of physical simulations, Merel et al. [48] presented a hierarchical RL framework that synthesizes movements related to catching and carrying boxes, utilizing egocentric observations. In a recent development, Hassan et al. [28] proposed a method based on adversarial imitation learning [53,81,57,5,56,19,98] for box manipulation. Furthermore, Xie et al. [84] showcased the manipulation of multiple boxes by integrating a trajectory planner with RL policy control. Although physics-based methods show promise in generalizing to the physical cues, in AR/VR or gaming applications, simplicity is valued. It is preferable to use a single model for synthesizing long-term, complex interactions rather than implementing complex methods with a family of motion policies. For this reason, our work aligns with kinematics-based approaches, aiming to provide a unified solution.

Among kinematics-based methods, the Neural State Machine (NSM) [67] stands out as a powerful method capable of modeling both static interactions, such as sitting, and dynamic interactions, including object manipulation. InterDiff [85] supports a more diverse set of objects by reasoning about interactions in the contact space. However, it primarily considers short-term interactions and assumes stable contact between human and objects. More recently, OMOMO [40] conditions on the object’s trajectory to predict full-body object manipulation, but it neglects the physical attributes of the interaction, such as the mass or friction. Our work leverages these overlooked physical attributes and synthesizes diverse human-object interactions with nuanced details.

Intuitive physics guided learning. The contemporary studies of intuitive physics [47,77,6] has experienced a resurgence. This is attributed to recent attempts to incorporate human-level intuitive physics capabilities into deep learn-




Fig. 2: The FORCE dataset accurately captures diverse, nuanced interaction motion under varying levels of resistance (N for Newtons). The last column shows results from the object shape augmentation (see Section 3), where nuanced motion details are preserved.

ing frameworks. Some work infers the intrinsic physical properties of the object such as mass or spring [20,96], while others predict the underlying object dynamics such as object trajectory or the forces [51,50,18,78,39,36,23], which we found to be more relevant. Some of the investigations in this field forecast object trajectories or forces using Newtonian physics principles from images [51,50], while others predict physical stability in moving object scenarios [78,39,36,23]. Applications of intuitive physics have recently expanded to 3D computer vision. These include recovering object trajectories by modeling gravity [18], and guiding 3D human pose estimation with center of mass and center of pressure.

We draw our key insight from the intuitive physics studies, and model intricate interplay between dynamic human forces and object’s resistance.

3 Dataset

Our goal is to build a model of human-object interactions that is aware of the intuitive physics, yet existing datasets fall short in accurately capturing human interactions with moving objects [67,25,92,8,34]. A recent dataset offers high-quality dynamic interaction data, but it does not capture different styled human motion under varying physical conditions [40] (see Table 1). To address this gap, we introduce the FORCE dataset, which comprises 450 motion sequences of human pushing, pulling and carrying objects. It includes accurately tracked human and objects. The dataset’s main characteristic is its diverse motion nuances through interactions with varying resistances. We capture interactions with 8

Table 1: Comparison with existing human-object interaction datasets with objects. The criteria include (i) high MoCap quality (ii) capture of moving objects (iii) RGB-D (iv) nuanced interaction with varying levels of resistance.

Dataset	High Quality	Moving Obj.	RGB-D	Resistance
NSM [67], SAMP [25]	✓	✗	✗	✗
COUCH [92]	✓	✗	✓	✗
BEHAVE [8], InterCap [34]	✗	✓	✓	✗
OMOMO [40]	✓	✓	✗	✗
Ours	✓	✓	✓	✓

objects at 3-6 levels of resistance. We hold out 30 sequences of data covering a range of interactions and scenarios for evaluation.

Human tracking. To overcome the susceptibility to noise and object occlusion, we improve upon existing depth-only camera-based capture systems [8,34,26]. We leverage a customized tracker, integrating 17 human-mounted IMU [54] sensors with the cameras. This enhancement improves the accuracy of the captured data. After initially fitting the SMPL[44] to the captured point clouds, we synchronize the tracking results with motion sequences recorded by IMU sensors. Subsequently, we optimize the body poses from the IMU tracking to align with the Kinect-fitted result. Further technical details about our human tracking methodology are available in the supplementary materials.

Objects tracking. We initiate by pre-scanning eight objects (carry: bin, suitcase, stool, basket, container, backpack; push/pull: chair, box) to obtain their template meshes. Subsequently, we fit them to annotated keypoints on the captured images using camera projection. The fitting is refined by running ICP to fit to the segmented object point cloud [13]. We consider our objects to be rigid.

Diverse interactions with varying resistance. After building our customized tracker, we systematically capture interactions with objects placed in varying positions relative to the human. To ensure the authenticity of the motion, the subject is only offered high-level guidance instead of specific instructions. Figure 2 showcases a selection of interactions. The capture is designed to cover diverse factors of variation. Firstly, we vary the resistance of the object in each interaction instance with removable weights. The same motion is performed for each object. Secondly, the dataset captures different contact modes including one-handed and two-handed interactions.

Note, our data collection is performed on a uniform surface, and we determine the frictional coefficient by measuring force meter readings during the uniform pulling of a 10kg object. For consistency, we approximate gravity (g) and the frictional coefficient (η) as 10N/kg and 1, respectively.

Object shape augmentation. To generalize to unseen shapes, each motion sequence is augmented with 10 differently shaped objects from one of the table, chair, bag, and shelf categories of ShapeNet [10] (see the last column of Figure 2). First, we detect the ground truth hand contact positions on the object surface based on a distance threshold of 0.05cm. Next, we transfer contacts by aligning the new object with the source via centering and scaling. The ground truth contacts are projected onto the nearest surface of the new object. At every

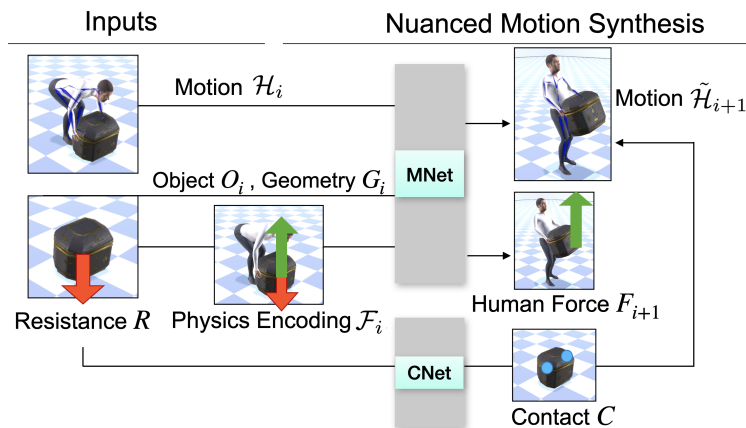


Fig. 3: Given the input human and object poses and the object geometry, our method which consists of *MNet* and *CNet*, synthesizes auto-repressively a diverse spectrum of nuanced interactions.

frame, following object sampling and rescaling, the human poses are recomputed to satisfy the contacts using inverse kinematics [67]. We assume the resistance remains unchanged after augmentation to preserve the context of human motion.

How does this dataset benefit the community?

- Developing physics-based synthesis methods [28,81,53].
- Synthesizing non-physics-informed object interactions [67,25,40,92,85].
- Human-object reconstruction [82,83,59,18].
- Benchmarking pose estimation for occlusion scenarios [14,63,15].
- Tracking human-object interaction with RGB-D [8,32].

4 Method

Our goal is to synthesize diverse, nuanced human-object interactions by modeling the physical attributes of resistance and the applied human force. The style of motion can be manipulated not only by varying the resistance, but also by the desired action and contact mode (e.g., left, right, two-handed).

4.1 Key Insight

Learning simultaneously from multi-modal human-object information, including the object pose, different modes of action and contact, can result in motions lacking in diversity and nuanced details. However, the intricate interplay between the force exerted by the human and the perceived resistance can help disambiguate between similar poses and facilitates learning from multi-class motions. There are two advantages to utilizing these two physical attributes. a) Human

force is a 3D representation that informs the model to which direction the force is applied, which dictates the human motion. b) Given the applied force, the resistance further refines the human motion. For instance, the human leans the body to compensate for the increased resistance.

Founded on this key insight, our method (Figure 3) utilizes a physics-aware motion network, *MNet* (Section 4.2) to auto-regressively synthesize diverse, nuanced interaction motion that is plausible to the physical conditions. More specifically, our method takes as input the initial human pose \mathcal{H}_0 , the object pose \mathbf{O}_0 and its geometry \mathbf{G}_0 , it auto-regressively predicts the future body motion $\mathcal{H}_{i>0}$. Notably, its physics awareness is achieved by our novel physics encoding. To ensure the motion is physically plausible, we use another neural network, *CNet* (Section 4.3) to predict the hand contact positions on the object surface.

4.2 MNet: Physics-aware Motion Prediction

At current frame i , given the poses of the human \mathcal{H}_i and the object \mathbf{O}_i , the object geometry \mathbf{O}_i , *MNet* leverages the interrelation between the applied human force and the perceived resistance to auto-regressively synthesize the future human motion \mathcal{H}_{i+1} . At the core of *MNet*, is the intuitive physics encoding \mathcal{F}_i , which encapsulates the crucial physics attributes of the interaction.

Intuitive Physics Encoding \mathcal{F}_i . Denoted as $\mathcal{F}_i = \{\mathbf{F}_i, \mathbf{R}, \mathbf{c}_i\}$, the encoding comprises the 3D human force $\mathbf{F}_i \in \mathbb{R}^{\mathcal{T} \times 3}$, the magnitude of resistance $\mathbf{R} \in \mathbb{R}$, and binary hand contact labels $C \in \{0, 1\}^{\mathcal{T} \times 2}$. The control signal is designed according to previous work [67,92,31,88,69], the signals are from the $\tau = 13$ uniformly sampled frames within the temporal window $[\max(0, i - 30), i + 30]$. This sampling ensures the model’s contextual awareness, allowing it to capture the temporal nuances of the interaction.

Next, we define the rest of the auto-regressive inputs:

Human motion \mathcal{H}_i : At frame i , the human motion comprises of the pose and the root trajectory: $\mathcal{H}_i = \{\mathbf{J}_i, \mathbf{T}_i\}$. The human pose $\mathbf{J}_i = (\mathbf{j}_i^p, \mathbf{j}_i^v, \mathbf{j}_i^r)$ contains root-relative joint positions $\mathbf{j}_i^p \in \mathbb{R}^{J \times 3}$, velocities $\mathbf{j}_i^v \in \mathbb{R}^{J \times 3}$, and rotations $\mathbf{j}_i^r \in \mathbb{R}^{J \times 6}$ (forward and upward vectors of the rotation matrix) for the $J = 22$ joints in SMPL skeleton [44].

The root trajectory $\mathbf{T}_i = (\mathbf{t}_i^p, \mathbf{t}_i^d, \mathbf{t}_i^a)$ encodes the root position $\mathbf{t}_i^p \in \mathbb{R}^{\mathcal{T} \times 2}$ and direction $\mathbf{t}_i^d \in \mathbb{R}^{\mathcal{T} \times 2}$ projected onto the ground. $\mathbf{t}_i^a \in [0, 1]^{\mathcal{T} \times 5}$ describes the current action (*idle, walk, carry, push, pull*). We use continuous values between 0 and 1 to facilitate the transition between actions.

Object \mathbf{O}_i : The target object is denoted as $\mathbf{O}_i = (\mathbf{o}_i^p, \mathbf{o}_i^d, \mathbf{o}_i^a)$, where $\mathbf{o}_i^p \in \mathbb{R}^{\mathcal{T} \times 3}$, $\mathbf{o}_i^d \in \mathbb{R}^{\mathcal{T} \times 6}$ are its positions and orientations relative to the root. $\mathbf{o}_i^a \in \{0, 1\}^{\mathcal{T} \times 5}$ are the binary variables describing the *desired future action*. For example, for carrying, the target action is “carry” when approaching the object.

Geometry \mathbf{G}_i : The object geometry \mathbf{G}_i is encoded by voxelizing object shape (in an $8 \times 8 \times 8$ dimensional grid). Each voxel stores its occupancy (\mathbb{R}) and the relative vector between the human root joint and the voxel (\mathbb{R}^3). This allows us to reason about the shape of the object and the distance between the person

and different parts of the object. We vectorize this grid to obtain our geometry encoding $\mathbf{G}_i \in \mathbb{R}^{2048}$.

Phase Φ_i : We introduce a learned 4-dimensional vector, phase $\Phi_i \in \mathbb{R}^4$, encoding the human joint trajectories following [66]. It encapsulates the spatial-temporal context of the motion.

Training. The *MNet* adopts a mixture-of-experts framework [31,25,67,68,92,91] and is trained in a supervised manner by directly minimizing the MSE loss on the following outputs.

$$\{\tilde{\mathcal{H}}_{i+1}, \mathbf{F}_{i+1}, \mathbf{O}_{i+1}, \Phi_{i+1}\} = f^{\text{MNet}}(\mathcal{H}_i, \mathcal{F}_i, \mathbf{O}_i, \mathbf{G}_i, \Phi_i), \quad (1)$$

where $\tilde{\mathcal{H}}_{i+1} = \{\mathbf{J}_{i+1}, \mathbf{T}_{i+1}, \tilde{\mathbf{j}}_{i+1}^p, \tilde{\mathbf{T}}_{i+1}\}$ denotes the future human motion. Here, \mathbf{J}_{i+1} , \mathbf{T}_{i+1} are the future body joints and the root trajectory. \mathbf{F}_{i+1} denotes the future 3D human forces. $\tilde{\mathbf{j}}_{i+1}^p$ and $\tilde{\mathbf{T}}_{i+1}$ are joint positions and root trajectory relative to the object. Supervising with these object-centric signals ensures the human reaches the target. \mathbf{O}_{i+1} and Φ_{i+1} are the object pose, and phase at frame $i + 1$. More details of the network architecture and training can be found in the supplementary.

Derivation of Human Force. For training, the human force is not directly observable. However, they can be estimated from the observable object acceleration. a) When the object moves, we leverage Newton’s laws to compute the human force. We assume the direction of force acts opposite to the direction of the object acceleration. For example, for *pushing* and *pulling*, the force is derived as $m \cdot g \cdot \eta \cdot \frac{\mathbf{a}}{\|\mathbf{a}\|} + m \cdot \mathbf{a}$, where m , g , η , \mathbf{a} denote the mass, gravity, frictional coefficient and object acceleration. b) When the object is in contact while remaining static, we derive the force by linearly interpolating between the zero, and the derived force from the nearest frame at which the object moves. Furthermore, to eliminate the effect of noise of object acceleration, we apply a Butterworth filter of kernel size 4, while preserving its range.

4.3 Resistance-conditioned Contact Prediction

Contacts are imperative to ensure the physical plausibility of the human-object interaction, and the way the human contacts with an object is influenced by its resistance. For example, when pushing a chair when the resistance is low, the human tends to push through the top of the chair since it is more efficient. However, when the resistance of the chair increases, the human tends to hold onto the lower part of the chair to stabilize to push more efficiently. Inspired by this and the previous work on predicting hand contact conditioned on the object geometry [92], we use a variational auto-encoder *CNet* to predict the hand contact positions, $\mathcal{C} \in \mathbb{R}^{2 \times 3}$, on the object surface. Formally, the contacts are predicted via $\hat{\mathcal{C}} = f^{\text{CNet}}(\mathbf{R}, \mathbf{G}_i, \mathbf{j}_i^p, \mathbf{o}_i^a)$. Here, \mathbf{R} , \mathbf{j}_i^p , \mathbf{o}_i^a are the resistive force, human joint positions and the target actions as defined previously. The object geometry \mathbf{G}_i in this case is encoded relative to the center of the object instead of the human’s root. During training, the network is trained to minimize the

following loss,

$$L_{\text{CNet}} = \|\hat{\mathcal{C}} - \mathcal{C}\|_2^2 + \lambda KL(q(\mathbf{z}|\bar{\mathbf{c}}, \mathbf{R}, \mathbf{G}, \mathbf{j}_i^p, \mathbf{o}_i^a)||p(\mathbf{z})), \quad (2)$$

where KL denotes the Kullback-Leibler divergence. For training, CNet is trained with ground truth contact positions derived from Section 3. For “one-handed” contact, the contact position for the other hand is zeroed. During inference, the contacts are sampled once the human is within the vicinity of the object, and is remained constant for the rest of the motion sequence. We apply a real-time inverse kinematics [25,67,92,91] to ensure the contact constraints are satisfied. We optimize the pose of the arm joints to minimize the distance between the linearly interpolated position between the current hand joint and contact. The weight of the interpolation is inversely proportional to the distance to ensure the smoothness and realism of the hand motion.

5 Experiments

For evaluation, we first introduce the baselines (Section 5.1). We then compare with the baselines quantitatively (Section 5.2) and qualitatively (Section 5.3). Last, we perform ablations on the designs of our method (Section 5.4) and show results on generalization (Section 5.5). *We encourage readers to refer to our supplementary video for animated qualitative results.*

5.1 Baselines

To our best knowledge, the most related works to ours are NSM [67] and SAMP [25] since they both synthesize human-object interactions. We train each baseline method on the FORCE dataset using their publicly released codes. These methods neglect the physics attributes of the interaction. To measure the effectiveness of our intuitive physics encoding, we adapt these baselines by additionally conditioning on the resistive force and the binary hand contact. We refer to these modified baselines as NSM+ and SAMP+. Moreover, we evaluate two ablative variations of FORCE. One model is without the physics encoding \mathcal{F} , and the other one is without the contact prediction.

5.2 Quantitative Evaluation

We quantitatively evaluate our model performance with two setups. a) We evaluate the accuracy (MPJPE) of motion synthesis on testing sequences following our training/testing split explained in Section 3. The model synthesizes poses auto-regressively given the trajectories. The sequences are each distinctive, covering diverse pushing, pulling and carrying motions at different resistances. b) We evaluate all other metrics on the complete synthesis of approaching and interacting in a random setup. For fairness, the object shapes are sampled. They are placed randomly relative to the human (within 6 meters), and given a random y-rotation. The synthesis is performed 120 times, repeating each action type (pushing, pulling, one-handed, and two-handed carrying) 30 times.

MPJPE(cm) [55,35]. The average per-joint error in centimeters, measuring the accuracy of human motion synthesis. Note, the result on “Ours no Contact” is omitted because the ablative baseline utilizes the same MNet as “Ours”.

Success rate(%) [11,25,28]. Evaluated on synthesis with randomly generated objects. An interaction is considered successful when the desired action and hand contact are both executed within 10 seconds.

Execution time(s) [27,53]. Average time in seconds till a successful interaction.

Diversity [24,87,93,25]. We compute the Average Pairwise Distance (APD) on the root-relative joint positions of the synthesized motion. Diversity is compared against the training data.

Collision Score(%) [93,90,27]. The percentage of frames that the body does not penetrate with the object.

Foot-slide(cm) [31,67,28,25,91,92]. The average distance traveled by the pivotal foot when in contact with the ground at each footstep.

Table 2: Quantitative comparisons with baselines. MPJPE is measured on the testing sequences, and the other metrics are evaluated on 120 syntheses performed with randomly sampled objects. Note, diversity on the dataset is *0.906*.

Method	MPJPE(cm) [↓]	Succ. Rate(%) [↑]	Exec. Time(s) [↓]	Diversity [→]	Collision(%) [↑]	Foot-slide(cm) [↓]
NSM [67]	7.71	75.1	5.94	0.707	64.3	4.59
SAMP [25]	8.67	66.2	7.72	0.719	57.7	5.79
NSM+	7.08	79.6	6.01	0.785	69.6	4.71
SAMP+	8.34	68.7	7.56	0.776	58.3	6.02
Ours no Phys. Enc. \mathcal{F}	6.75	86.7	5.81	0.847	75.4	4.60
Ours no Contact Pred.	-	94.2	5.50	0.886	77.2	4.48
Ours	6.02	97.5	5.42	0.891	84.0	4.46

From Table 2, it can be seen that FORCE attains the lowest MPJPE of 6.02 cm, surpassing both baseline models and ablative variations of FORCE. This emphasizes the crucial role of our intuitive physics encoding in facilitating the learning of multi-class motions.

Under the simulated object setup, FORCE has the highest success rate of interaction of 97.5% and the shortest average execution time of 5.42 seconds, illustrating the better controllability of our approach.

Our model synthesizes motion of diversity (0.891) which is the closest match to the training data distribution (0.906). This underscores the significance of the novel physics encoding in enabling the synthesis of diverse, nuanced interactions. Additionally, our method minimizes object interpenetration, as reflected in the highest collision scores. It is notable that, compared with the ablative baseline without contact prediction, our method has a higher collision score. This indicates that with contact prediction, our model synthesizes more physically plausible motion. Furthermore, our model demonstrates improved motion realism, evident in its reduced foot-slide.

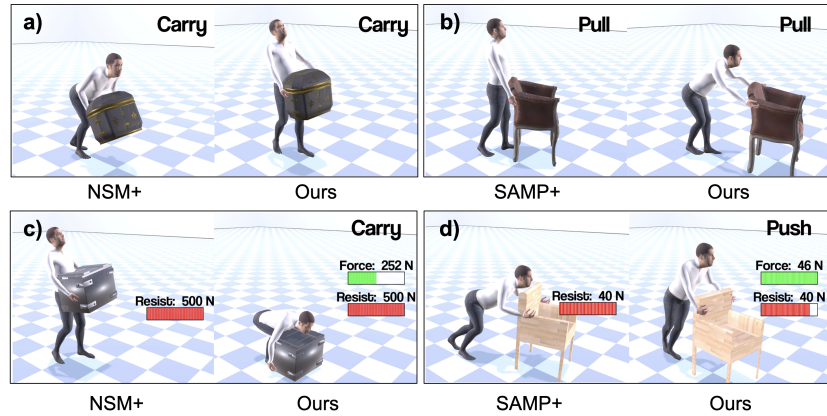


Fig. 4: Comparison with baselines. Our *intuitive physics encoding* \mathcal{F} disambiguates different multi-class motions and enables the synthesis of the desired interaction (carrying in **a**), pulling in **b**). **c**) and **d**) show the coupled encoding of human force and resistance allows the synthesis of visually plausible interactions. In **c**), when the resistance is greater than the human force, FORCE synthesizes a more plausible “infeasible” interaction. In **d**), with low resistance, and the human pose maintains an upright position.

5.3 Qualitative Evaluation

Figure 4 illustrates our qualitative comparison with the baselines NSM+ [67] and SAMP+ [25]. In Figure 4 a), the poses when lifting up an object and pushing an object are similar, introducing ambiguity between these two classes of motions. However, the physics encoding informs the model about the crucial directional information of the human force, dictating the model to synthesize the desired carrying motion. Similarly in Figure 4 b), our model synthesizes the desired pulling motion. The physics encoding also provides awareness of the feasibility given the resistance, because of the coupled encoding of human force and resistances. In Figure 4 c), when the resistance is greater than the exerted human force, FORCE is aware that the interaction is no longer feasible. In Figure 4 d), our synthesis is more visually plausible given the low resistance, with the body pose remaining upright.

5.4 Ablations

Physics encoding \mathcal{F} . As shown in Figure 5, FORCE synthesizes nuanced human-object interactions with objects of the same shape but varying resistances. Empowered by the intuitive physics encoding \mathcal{F} , our method synthesizes diverse interactions. It generalizes not only to different resistances of the object but also to different actions and contacts. The graph on the right provides a zoomed-in analysis of the motion, where we evaluate the horizontal shift of the

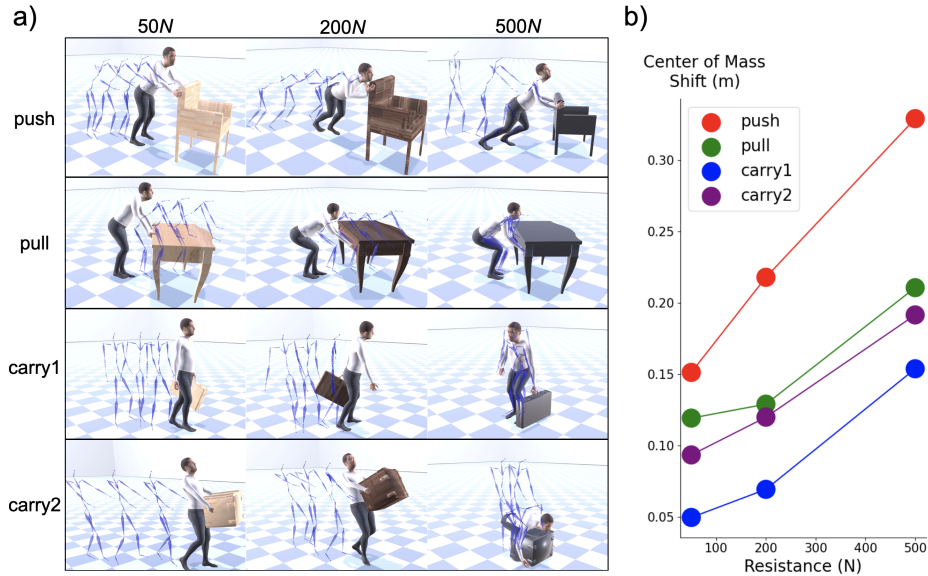


Fig. 5: Nuanced interactions with the same object shape, but different resistance. a) The human struggles with the interaction as the resistance increases from 50 N to 200 N. At 500 N, the human fails to interact with the object. b) Plot depicts the horizontal shift of the center of mass of the human, estimated following [73,12]. Results averaged over 1000 frames.




Fig. 6: Without Physics Encoding (left) vs. FORCE (right). With physics encoding, synthesizes visually plausible motion, when the resistance is greater than the applied human force.




Fig. 7: Without geometry (left) vs. FORCE (right). With geometry representation \mathbf{G} , reasons about collision and to avoid interpenetration.

center of mass of the human (for each 1000-frame sequence). The center of mass of the human is estimated following [73,12]. It can be observed from the plot, that there is a direct positive correlation between the perceived resistance of the human during the interaction and the shift of center of mass.

Figure 6 also emphasizes the importance of intuitive physics encoding. On the left, without the physics encoding, the model overlooks the resistance surpassing the applied force of the human. With the physics encoding on the right, our

method is aware that the resistance is not feasible, as the perceived resistance reaches 500 N and is greater than the applied human force.

Geometry representation G. With the representation introduced in Section 4.2, our method reasons about the collision as the voxels provide coarse geometry details. As a result, it exhibits less interpenetration (Figure 7).

Contact prediction. As shown Figure 8 a), different contact positions are predicted by conditioning on different resistances using our method. It can be concluded that the contact depends on the resistance. Here, as the resistance increases, the human tends to push from the lower part of the chair to push more efficiently. Figure 8 b) shows, with the contact prediction, FORCE synthesizes motion that satisfies better the hand contact constraints on the surface, and exhibits less interpenetration (see Table 2).




Fig. 8: a) Different hand contacts predicted by FORCE. Contact is dependent on the resistance. b) Without contact prediction (left) vs. FORCE (right). FORCE synthesizes more physically plausible motion.

5.5 Generalization to different shapes and locations.

As illustrated in Figure 9, FORCE generalizes to different shapes. Thanks to the object shape augmentation introduced in Section 3, our model is able to synthesize lifting a chair, an object category that is not seen during training of carrying sequences.



Fig. 9: (Top) FORCE generalizes to unseen shapes, demonstrated via pushing, carrying, and pulling. (Bottom) FORCE generalizes to different locations. one-handed and two-handed carry from the table and from the floor.

It is also highlighted in Figure 9 that FORCE interacts with objects at different locations, for example, picking up from different heights.

6 Conclusion

This work, FORCE, tackles the problem of synthesizing human object interactions with nuanced details, by modeling intuitive physics. It solves two main challenges. First, it proposes the first kinematic-based method that synthesizes human-object interaction conditioned on physical attributes such as object resistance and the human force. By leveraging the novel physics encoding, our method generates a diverse spectrum of interaction motions under varying resistance. Its performance in diversity and realism surpasses previous methods. Second, the accompanying FORCE dataset, featuring over 450 motion sequences, provides a valuable resource for training and evaluating kinematic-based methods.

For limitations and future work, this work paved the way for research in nuanced human-object interaction in more complex scenarios. First, while our model demonstrates its efficacy in generating diverse human-object interactions, there is an opportunity for expansion by considering a broader range of subjects, because each individual possesses unique strengths and approaches to interactions. Second, interaction can be extended to complex scenarios of dynamic resistance. For example, when carrying a tank of water, the resistance fluctuates and influences human motion. We release the code, dataset, and models, to stimulate future research.

Acknowledgement

We thank RVH group members for their helpful discussions. This work is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 409792180 (Emmy Noether Programme, project: Real Virtual Humans), and German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A, and Huawei Noah’s Ark Lab. Gerard Pons-Moll is a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645. The project was made possible by funding from the Carl Zeiss Foundation.

References

1. Microsoft azure kinect. https://en.wikipedia.org/wiki/Azure_Kinect/
2. Alexanderson, S., Nagy, R., Beskow, J., Henter, G.E.: Listen, denoise, action! audio-driven motion synthesis with diffusion models. *ACM Trans. Graph.* **42**(4) (2023)
3. Alldieck, T., Magnor, M., Bhatnagar, B.L., Theobalt, C., Pons-Moll, G.: Learning to reconstruct people in clothing from a single RGB camera. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (jun 2019)
4. Araújo, J.P., Li, J., Vetrivel, K., Agarwal, R., Wu, J., Gopinath, D., Clegg, A., Liu, C.K.: CIRCLE: capture in rich contextual environments. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR* (2023)
5. Bae, J., Won, J., Lim, D., Min, C., Kim, Y.M.: PMP: learning to physically interact with environments using part-wise motion priors. In: Brunvand, E., Sheffer, A., Wimmer, M. (eds.) *ACM SIGGRAPH 2023 Conference Proceedings, SIGGRAPH 2023*
6. Battaglia, P.W., Pascanu, R., Lai, M., Rezende, D.J., Kavukcuoglu, K.: Interaction networks for learning about objects, relations and physics. In: *Advances in Neural Information Processing Systems 2016*
7. Bhatnagar, B.L., Sminchisescu, C., Theobalt, C., Pons-Moll, G.: Combining implicit function learning and parametric models for 3d human reconstruction. In: *European Conference on Computer Vision (ECCV)*. Springer (August 2020)
8. Bhatnagar, B.L., Xie, X., Petrov, I., Sminchisescu, C., Theobalt, C., Pons-Moll, G.: Behave: Dataset and method for tracking human object interactions. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2022)
9. Braun, J., Christen, S., Kocabas, M., Aksan, E., Hilliges, O.: Physically plausible full-body hand-object interaction synthesis
10. Chang, A.X., Funkhouser, T.A., Guibas, L.J., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: Shapenet: An information-rich 3d model repository. *CoRR* **abs/1512.03012** (2015)
11. Chao, Y., Yang, J., Chen, W., Deng, J.: Learning to sit: Synthesizing human-chair interactions via hierarchical control. *CoRR* **abs/1908.07423** (2019)
12. Chebel, E., Tunc, B.: Evaluation of center of mass estimation for obese using statically equivalent serial chain. *Scientific Report* **12** (2022)
13. Cheng, H.K., Schwing, A.G.: XMem: Long-term video object segmentation with an atkinson-shiffrin memory model. In: *ECCV* (2022)
14. Cheng, Y., Yang, B., Wang, B., Tan, R.T.: 3d human pose estimation using spatio-temporal networks with explicit occlusion training. In: *The Conference on Artificial Intelligence AAAI 2020*
15. Cheng, Y., Yang, B., Wending, Y., Tan, R.T.: Occlusion-aware networks for 3d human pose estimation in video. In: *International Conference on Computer Vision, ICCV* (2019)
16. Christen, S., Kocabas, M., Aksan, E., Hwangbo, J., Song, J., Hilliges, O.: D-grasp: Physically plausible dynamic grasp synthesis for hand-object interactions. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022)
17. Dabral, R., Mughal, M.H., Golyanik, V., Theobalt, C.: Mofusion: A framework for denoising-diffusion-based motion synthesis. In: *Computer Vision and Pattern Recognition (CVPR)* (2023)
18. Dabral, R., Shimada, S., Jain, A., Theobalt, C., Golyanik, V.: Gravity-aware monocular 3d human-object reconstruction. In: *International Conference on Computer Vision (ICCV)* (2021)

19. Dou, Z., Chen, X., Fan, Q., Komura, T., Wang, W.: Ase: Learning conditional adversarial skill embeddings for physics-based characters. CoRR **abs/2309.11351** (2023)
20. Duan, J., Dasgupta, A., Fischer, J., Tan, C.: A survey on machine learning approaches for modelling intuitive physics. In: Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022
21. Eigen, D., Ranzato, M., Sutskever, I.: Learning factored representations in a deep mixture of experts. In: Bengio, Y., LeCun, Y. (eds.) 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings (2014)
22. Ghosh, A., Dabral, R., Golyanik, V., Theobalt, C., Slusallek, P.: Imos: Intent-driven full-body motion synthesis for human-object interactions. In: Eurographics (2023)
23. Groth, O., Fuchs, F.B., Posner, I., Vedaldi, A.: Shapestacks: Learning vision-based physical intuition for generalised object stacking. In: European Conference on Computer Vision - ECCV 2018
24. Guo, C., Zou, S., Zuo, X., Wang, S., Ji, W., Li, X., Cheng, L.: Generating diverse and natural 3d human motions from text. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
25. Hassan, M., Ceylan, D., Villegas, R., Saito, J., Yang, J., Zhou, Y., Black, M.: Stochastic scene-aware motion prediction. In: Proceedings of the International Conference on Computer Vision 2021 (Oct 2021)
26. Hassan, M., Choutas, V., Tzionas, D., Black, M.J.: Resolving 3D human pose ambiguities with 3D scene constraints. In: International Conference on Computer Vision (Oct 2019)
27. Hassan, M., Ghosh, P., Tesch, J., Tzionas, D., Black, M.J.: Populating 3D scenes by learning human-scene interaction. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2021) (2021)
28. Hassan, M., Guo, Y., Wang, T., Black, M.J., Fidler, S., Peng, X.B.: Synthesizing physical character-scene interactions. CoRR **abs/2302.00883** (2023)
29. Henter, G.E., Alexanderson, S., Beskow, J.: Moglow: probabilistic and controllable motion synthesis using normalising flows. ACM Trans. Graph. **39**(6) (2020)
30. Holden, D., Kanoun, O., Perepichka, M., Popa, T.: Learned motion matching. ACM Trans. Graph. **39**(4) (2020)
31. Holden, D., Komura, T., Saito, J.: Phase-functioned neural networks for character control. ACM Trans. Graph. **36**(4) (2017)
32. Huang, C.H.P., Yi, H., Höschle, M., Safroshkin, M., Alexiadis, T., Polikovsky, S., Scharstein, D., Black, M.J.: Capturing and inferring dense full-body human-scene contact. In: Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR) (Jun 2022)
33. Huang, S., Wang, Z., Li, P., Jia, B., Liu, T., Zhu, Y., Liang, W., Zhu, S.C.: Diffusion-based generation, optimization, and planning in 3d scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
34. Huang, Y., Taheri, O., Black, M.J., Tzionas, D.: InterCap: Joint markerless 3D tracking of humans and objects in interaction. In: German Conference on Pattern Recognition (GCPR) (2022)
35. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. IEEE Transactions on Pattern Analysis and Machine Intelligence (2014)

36. Janner, M., Levine, S., Freeman, W.T., Tenenbaum, J.B., Finn, C., Wu, J.: Reasoning About Physical Interactions with Object-Oriented Prediction and Planning. In: International Conference on Learning Representations, ICLR 2019
37. Kulkarni, N., Rempe, D., Genova, K., Kundu, A., Johnson, J., Fouhey, D., Guibas, L.: Nifty: Neural object interaction fields for guided human motion synthesis. arXiv:2307.07511 (2023)
38. Lee, J., Joo, H.: Locomotion-action-manipulation: Synthesizing human-scene interactions in complex 3d environments. arXiv preprint arXiv:2301.02667 (2023)
39. Lerer, A., Gross, S., Fergus, R.: Learning physical intuition of block towers by example. In: International Conference on Machine Learning, ICML 2016
40. Li, J., Wu, J., Liu, C.K.: Object motion guided human motion synthesis. ACM Transactions on Graphics **42**(6) (Dec 2023)
41. Li, Q., Wang, J., Loy, C.C., Dai, B.: Task-oriented human-object interactions generation with implicit neural representations (2023)
42. Li, R., Yang, S., Ross, D.A., Kanazawa, A.: Ai choreographer: Music conditioned 3d dance generation with aist++ (2021)
43. Li, X., Liu, S., Kim, K., Wang, X., Yang, M.H., Kautz, J.: Putting humans in a scene: Learning affordance in 3d indoor environments. In: IEEE Conference on Computer Vision and Pattern Recognition (2019)
44. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. ACM Trans. Graphics (Proc. SIGGRAPH Asia) **34** (Oct 2015)
45. Mao, W., Liu, M., Salzmann, M., Li, H.: Learning trajectory dependencies for human motion prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2019)
46. Martinez, J., Black, M.J., Romero, J.: On human motion prediction using recurrent neural networks. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017
47. McCloskey, M., Washburn, A., Felch, L.: Intuitive physics: The straight-down belief and its origin. Journal of experimental psychology. Learning, memory, and cognition **9**, 636–49 (10 1983)
48. Merel, J., Tunyasuvunakool, S., Ahuja, A., Tassa, Y., Hasenclever, L., Pham, V., Erez, T., Wayne, G., Heess, N.: Catch & carry: reusable neural controllers for vision-guided whole-body tasks. ACM Trans. Graph. (2020)
49. Mir, A., Puig, X., Kanazawa, A., Pons-Moll, G.: Generating continual human motion in diverse 3d scenes. In: International Conference on 3D Vision (3DV) (March 2024)
50. Mottaghi, R., Bagherinezhad, H., Rastegari, M., Farhadi, A.: Newtonian image understanding: Unfolding the dynamics of objects in static images. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016
51. Mottaghi, R., Rastegari, M., Gupta, A., Farhadi, A.: "what happens if..." learning to predict the effect of forces in images. In: European Conference on Computer Vision - ECCV 2016
52. Pan, L., jingbo Wang, Huang, B., Zhang, J., Wang, H., Tang, X., Wang, Y.: Synthesizing physically plausible human motions in 3d scenes (2023)
53. Pan, L., Wang, J., Huang, B., Zhang, J., Wang, H., Tang, X., Wang, Y.: Synthesizing physically plausible human motions in 3d scenes (2023)
54. Paulich, M., Schepers, M., Rudigkeit, N., Bellusci, G.: Xsens mtw awinda: Miniature wireless inertial magnetic motion tracker for highly accurate 3d kinematic applications (2018)

55. Pavlo, D., Grangier, D., Auli, M.: Quaternet: A quaternion-based recurrent model for human motion. In: British Machine Vision Conference (BMVC) (2018)
56. Peng, X.B., Guo, Y., Halper, L., Levine, S., Fidler, S.: Ase: Large-scale reusable adversarial skill embeddings for physically simulated characters. *ACM Trans. Graph.* (2022)
57. Peng, X.B., Ma, Z., Abbeel, P., Levine, S., Kanazawa, A.: AMP: adversarial motion priors for stylized physics-based character control. *ACM Trans. Graph.* (2021)
58. Pérez, G.V., Henter, G.E., Beskow, J., Holzapfel, A., Oudeyer, P., Alexanderson, S.: Transflower: probabilistic autoregressive dance generation with multimodal attention. *ACM Trans. Graph.* **40**(6) (2021)
59. Petrov, I.A., Marin, R., Chibane, J., Pons-Moll, G.: Object pop-up: Can we infer 3d objects and their poses from human interactions alone? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2023)
60. Pi, H., Peng, S., Yang, M., Zhou, X., Bao, H.: Hierarchical generation of human-object interactions with diffusion probabilistic models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2023)
61. Qi, S., Zhu, Y., Huang, S., Jiang, C., Zhu, S.C.: Human-centric indoor scene synthesis using stochastic grammar. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
62. Rong, Y., Shiratori, T., Joo, H.: Frankmocap: A monocular 3d whole-body pose estimation system via regression and integration. In: IEEE International Conference on Computer Vision Workshops (2021)
63. Sáráandi, I., Linder, T., Arras, K.O., Leibe, B.: How robust is 3d human pose estimation to occlusion? <http://arxiv.org/abs/1808.09316>
64. Shafir, Y., Tevet, G., Kapon, R., Bermano, A.H.: Human motion diffusion as a generative prior. arXiv preprint arXiv:2303.01418 (2023)
65. Shi, Y., Wang, J., Jiang, X., Dai, B.: Controllable motion diffusion model. CoRR **abs/2306.00416**, <https://doi.org/10.48550/arXiv.2306.00416>
66. Starke, S., Mason, I., Komura, T.: Deepphase: periodic autoencoders for learning motion phase manifolds. *ACM Trans. Graph.* (2022)
67. Starke, S., Zhang, H., Komura, T., Saito, J.: Neural state machine for character-scene interactions. *ACM Trans. Graph.* **38**(6) (2019)
68. Starke, S., zhao, Y., Komura, T., Zaman, K.A.: Local motion phases for learning multi-contact character movements. *ACM Trans. Graph.*
69. Starke, S., zhao, Y., Zinno, F., Komura, T.: Neural animation layering for synthesizing martial arts movements. *ACM Trans. Graph.*
70. Taheri, O., Choutas, V., Black, M.J., Tzionas, D.: GOAL: Generating 4D whole-body motion for hand-object grasping. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
71. Taheri, O., Zhou, Y., Tzionas, D., Zhou, Y., Ceylan, D., Pirk, S., Black, M.J.: Grip: Generating interaction poses conditioned on object and body motion. In: International Conference on 3D Vision (3DV 2024) (2024)
72. Tevet, G., Raab, S., Gordon, B., Shafir, Y., Cohen-or, D., Bermano, A.H.: Human motion diffusion model. In: The Eleventh International Conference on Learning Representations (2023)
73. Tripathi, S., Muller, L., Huang, C.H.P., Omid, T., Black, M.J., Tzionas, D.: 3d human pose estimation via intuitive physics. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
74. Wang, J., Xu, H., Xu, J., Liu, S., Wang, X.: Synthesizing long-term 3d human motion and interaction in 3d scenes. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2021)

75. Wang, J., Rong, Y., Liu, J., Yan, S., Lin, D., Dai, B.: Towards diverse and natural scene-aware 3d human motion synthesis. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. IEEE
76. Wang, Z., Chen, Y., Liu, T., Zhu, Y., Liang, W., Huang, S.: Humanise: Language-conditioned human motion generation in 3d scenes. In: Advances in Neural Information Processing Systems (NeurIPS) (2022)
77. Watters, N., Zoran, D., Weber, T., Battaglia, P.W., Pascanu, R., Tacchetti, A.: Visual interaction networks: Learning a physics simulator from video. In: Advances in Neural Information Processing Systems 2017
78. Wu, J., Lu, E., Kohli, P., Freeman, B., Tenenbaum, J.: Learning to see physics via visual de-animation. In: Advances in Neural Information Processing Systems 2017
79. Wu, Y., Wang, J., Zhang, Y., Zhang, S., Hilliges, O., Yu, F., Tang, S.: Saga: Stochastic whole-body grasping with contact. In: Proceedings of the European Conference on Computer Vision (ECCV) (2022)
80. Wu, Y., Kirillov, A., Massa, F., Lo, W.Y., Girshick, R.: Detectron2. <https://github.com/facebookresearch/detectron2> (2019)
81. Xiao, Z., Wang, T., Wang, J., Cao, J., Dai, B., Lin, D., Pang, J.: Unified human-scene interaction via prompted chain-of-contacts. Arxiv (2023)
82. Xie, X., Bhatnagar, B.L., Pons-Moll, G.: Chore: Contact, human and object reconstruction from a single rgb image. In: European Conference on Computer Vision (ECCV). Springer (October 2022)
83. Xie, X., Bhatnagar, B.L., Pons-Moll, G.: Visibility aware human-object interaction tracking from single rgb camera. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2023)
84. Xie, Z., Tseng, J., Starke, S., van de Panne, M., Liu, C.K.: Hierarchical planning and control for box loco-manipulation. CoRR **abs/2306.09532**
85. Xu, S., Li, Z., Wang, Y.X., Gui, L.Y.: Interdiff: Generating 3d human-object interactions with physics-informed diffusion. In: ICCV (2023)
86. Xuan, H., Li, X., Jinsong Zhang, H.Z., Liu, Y., Li, K.: Narrator: Towards natural control of human-scene interaction generation via relationship reasoning. In: ICCV (2023)
87. Yuan, Y., Kitani, K.: Dlow: Diversifying latent flows for diverse human motion prediction. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
88. Zhang, H., Starke, S., Komura, T., Saito, J.: Mode-adaptive neural networks for quadruped motion control. ACM Trans. Graph. **37**(4) (2018)
89. Zhang, H., Ye, Y., Shiratori, T., Komura, T.: Manipnet: neural manipulation synthesis with a hand-object spatial representation. ACM Trans. Graph. (2021)
90. Zhang, S., Zhang, Y., Ma, Q., Black, M.J., Tang, S.: PLACE: Proximity learning of articulation and contact in 3D environments. In: International Conference on 3D Vision (3DV) (Nov 2020)
91. Zhang, W., Dabral, R., Leimkühler, T., Golyanik, V., Habermann, M., Theobalt, C.: Roam: Robust and object-aware motion generation using neural pose descriptors. arXiv preprint arXiv:2308.12969 (2023)
92. Zhang, X., Bhatnagar, B.L., Starke, S., Guzov, V., Pons-Moll, G.: Couch: Towards controllable human-chair interactions. In: European Conference on Computer Vision (ECCV) (October 2022)
93. Zhang, Y., Hassan, M., Neumann, H., Black, M.J., Tang, S.: Generating 3d people in scenes without people. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)

94. zhao, K., Wang, S., Zhang, Y., Beeler, T., , Tang, S.: Compositional human-scene interaction synthesis with semantic control. In: European conference on computer vision (ECCV) (2022)
95. Zhao, K., Zhang, Y., Wang, S., Beeler, T., , Tang, S.: Synthesizing diverse human motions in 3d indoor scenes. In: International conference on computer vision (ICCV) (2023)
96. Zheng, Q., Wu, W., Pan, H., Mitra, N.J., Cohen-Or, D., Huang, H.: Inferring object properties from human interaction and transferring them to new motions (2021)
97. Zhou, K., Bhatnagar, B.L., Lenssen, J.E., Pons-Moll, G.: Toch: Spatio-temporal object-to-hand correspondence for motion refinement. In: European Conference on Computer Vision (ECCV) (2022)
98. Zhu, Q., Zhang, H., Lan, M., Han, L.: Neural categorical priors for physics-based character control (2023)

APPENDIX

1 List of Symbols

Acknowledging the abundance of symbols in this paper, we aim to clarify and eliminate any confusion by presenting a detailed list of symbols (referenced in Table 6) used in this paper, along with their corresponding definitions.

Notation	Definition
\mathcal{F}	intuitive physics encoding
\mathbf{R}	magnitude of the resistance
\mathbf{F}	3D human force
\mathbf{c}	binary desired hand contact
\mathcal{H}	human motion
\mathbf{J}	human pose
\mathbf{j}_i^p	joint position (root-relative)
\mathbf{j}_i^v	joint velocity (root-relative)
\mathbf{j}_i^r	joint rotation (root-relative)
\mathbf{T}	human trajectory
\mathbf{t}_i^p	ground-projected root position
\mathbf{t}_i^d	ground-projected root rotation
\mathbf{t}_i^a	current action
\mathbf{O}	object
\mathbf{o}_i^p	object position (root-relative)
\mathbf{o}_i^d	object rotation (root-relative)
\mathbf{o}_i^a	binary desired action
\mathbf{G}	geometry representation computed from the object voxels
ϕ	phase encoding of the human motion
\mathcal{C}	hand contact positions

Table 3: Legend of Notations.

2 Dataset

Further details on the dataset: Our dataset comprises 450 motion sequences involving human-object interactions with a diverse range of resistance forces. Table 4 provides a detailed breakdown of the dataset categorized by the level of resistance. In this context, resistance is measured solely by the mass of the removable weight used. The masses of the objects themselves are not factored into these measurements. They are measured and will be provided with the dataset. The dataset distributions based on the type of action (Table 5), and the type of hand contact (Table 6) are also presented.

Table 4: Distribution of the dataset by the level of resistance. The data is categorized by the mass of removable weight used. Note, the masses of the objects themselves are not factored into these measurements.

Mass	Minutes	%
0 kg	47.1	33.3
5 kg	18.2	12.9
10 kg	21.4	15.1
15 kg	23.8	16.8
20 kg	7.9	5.6
25 kg	8.9	6.2
>30 kg	14.2	10.1

Table 5: Distribution of the dataset by the type of action.

Action Type	Minutes	%
Carry	107.3	75.9
Push	17.4	12.3
Pull	16.7	11.8

Table 6: Distribution of the dataset with different hand contact.

Interaction Type	Minutes	%
Right Hand	22.7	19.4
Left Hand	27.4	16.0
Both Hand	91.4	64.6

Details on human tracking. The first stage of our human tracking is to fit the SMPL parametric model [44] to the point clouds captured by the Kinect cameras. We segment humans in captured RGB images using Detectron V2 [80]. The resulting masks are then used to segment the human from the RGB data, before the the human point cloud is lifted in 3D. To initialize the SMPL pose, we employ FrankMocap [62] from the images. Subsequently, instance-specific optimization techniques [3] are applied to fit the SMPL model to the segmented human point cloud via ICP. For more precise fitting, we further derive the SMPL shape parameters of each subject from 3D scans using [7]. This stage produces the SMPL parameters fitted to the cameras, but they can be noisy and erroneous due to occlusion.

The second stage of our tracking is to refine the IMU-captured motion, which is smoother and more robust against occlusion. We synchronize the IMU-captured motion with the Kinect-fitted results from the previous stage, then perform an optimization to further refine the IMU-captured motion with the previously fitted results. The resulting motion is smooth and accurately captures the contact between the human and the object.

3 Architecture and Training Details

The motion synthesis network, *MNet*, adopts a mixture-of-expert structure [21]. Both the gating network and the prediction networks consist of three-layer fully-connected networks, with hidden dimensions of 128 and 512, respectively. The model employs 8 experts and is trained for 150 epochs using an Adam optimizer.

The initial learning rate is set at 1e-4, and a cosine learning rate scheduler gradually reduces it to 5e-6. A batch size of 32 is utilized, and the complete training process takes approximately 9 hours on an NVIDIA V100 GPU.

The contact prediction network, *CNet* encodes the object geometry \mathbf{G} through a three-layer fully connected network of shape {512, 512, 64}, the resistance \mathbf{R} , human joint positions \mathbf{j}_i^p and desired action \mathbf{o}_i^a in a separate network with identical shape. The latent vector \mathbf{z} of the VAE is of size 6. The weight of the Kullback-Leibler divergence β is 0.1. We use the Adam optimizer with a learning rate of 1e-3 and train CNet for 150 epochs. The full training of a subject-specific model takes approximately 10 minutes on an NVIDIA V100 GPU.

4 Ethics, Bias, and Boarder Impact

Our dataset was captured with the full consent and cooperation of the participant, ensuring ethical considerations were prioritized throughout the process. To mitigate potential biases, deliberate steps were taken. First, we captured objects of a range of shapes and sizes, and performed object shape augmentation to avoid biases arising from over- or under-representation of certain object shapes. Secondly, we carefully considered the distribution of resistance levels to prevent bias towards specific interaction intensities, ensuring a balanced representation of challenging and less challenging scenarios. Environmental factors, such as floor surface and lighting conditions, were also meticulously controlled and maintained constant throughout the data capture process to minimize any confounding variables.

Our dataset addresses the scarcity of accurate 3D human-object interaction data. It is the first dataset to capture diverse motions with detailed nuances, considering the physical aspects of the interaction. A large-scale 3D human motion dataset holds significant potential to catalyze novel research across various fields, including animation, gaming, virtual reality, and human-computer interaction. As of now, human motion datasets and synthesis methods have not resulted in any negative social impact.