ARCTIC: A Dataset for Dexterous Bimanual Hand-Object Manipulation

Zicong Fan[1,3] Omid Taheri[3] Dimitrios Tzionas[2] Muhammed Kocabas[1,3]
Manuel Kaufmann[1] Michael J. Black[3] Otmar Hilliges[1]

[1]ETH Zürich, Switzerland [2]University of Amsterdam [3]Max Planck Institute for Intelligent Systems, Tübingen, Germany

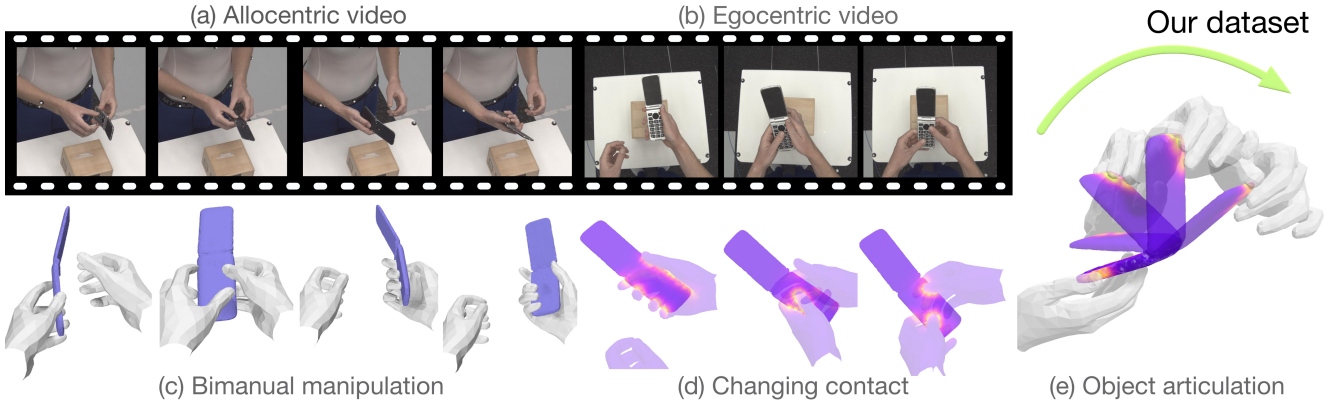


Figure 1. ARCTIC is a dataset of hands dexterously manipulating *articulated* objects. The dataset contains videos from both eight 3^{rd}-person allocentric views (a) and one 1^{st}-person egocentric view (b), together with accurate ground-truth 3D hand and object meshes, captured with a high-quality motion capture system. ARCTIC goes beyond existing datasets to enable the study of dexterous bimanual manipulation of articulated objects (c) and provides detailed contact information between the hands and objects during manipulation (d-e).

Abstract

Humans intuitively understand that inanimate objects do not move by themselves, but that state changes are typically caused by human manipulation (e.g., *the opening of a book). This is not yet the case for machines. In part this is because there exist no datasets with ground-truth 3D annotations for the study of physically consistent and synchronised motion of hands and articulated objects. To this end, we introduce ARCTIC – a dataset of two hands that dexterously manipulate objects, containing* 2.1*M video frames paired with accurate 3D hand and object meshes and detailed, dynamic contact information. It contains bi-manual articulation of objects such as scissors or laptops, where hand poses and object states evolve jointly in time. We propose two novel articulated hand-object interaction tasks: (1) Consistent motion reconstruction: Given a monocular video, the goal is to reconstruct two hands and articulated objects in 3D, so that their motions are spatio-temporally consistent. (2) Interaction field estimation: Dense relative hand-object distances must be estimated from images. We introduce two baselines ArcticNet and InterField, respectively and evaluate them qualitatively and quantitatively on ARCTIC. Our code and data are available at* https://arctic.is.tue.mpg.de.

1. Introduction

Humans constantly manipulate complex objects: we open our laptop's cover to work, we apply spray to clean, we carefully control our fingers to cut with scissors – rigid and articulated parts of objects move *together* with our hands. Inanimate objects only move or deform if external forces are applied to them. The study of the physically consistent dynamics of hands and objects during manipulation has so far been under-researched in the hand pose estimation literature. This is partly because existing hand-object datasets [8, 22, 23, 25, 37, 41] are mostly limited to grasping of rigid objects and contain few if any examples of rich and dexterous manipulation of articulated objects.

To enable the study of dexterous articulated hand-object manipulation, we collect a novel dataset called ARCTIC (**AR**ticulated obje**CT**s in **I**ntera**C**tion). ARCTIC consists of video sequences of multi-view RGB frames, and each frame is paired with accurate 3D hand and object meshes. ARCTIC contains data from 10 subjects interacting with 11 articulated objects, resulting in a total of 2.1M RGB images. Images are captured from multiple synchronized and calibrated views, including 8 static allocentric views and 1 moving egocentric view. To capture accurate 3D meshes

during manipulation, we synchronize color cameras with 54 high-resolution Vicon MoCap cameras [75]. These allow the use of small MoCap markers that do not interfere with hand-object interaction and are barely visible in the images. We then fit pre-scanned human and object meshes to the observed markers [42, 65]. The objects consist of two rigid parts that rotate about a shared axis such as the flip phone in Fig. 1 (for all objects, see SupMat).

Our dataset enables two novel tasks: (1) consistent motion reconstruction, (2) interaction field estimation. For *consistent motion reconstruction*, given a monocular video, the task is to reconstruct the 3D motion of two hands and an articulated object. In particular, the reconstructed hand-object meshes should have spatio-temporally consistent hand-object contact, object articulation, and smooth motion during interaction. This task has several challenges: (1) Spatio-temporal consistency requires precise hand-object 3D alignment for all frames; (2) This precision is hard to achieve due to depth ambiguity and severe occlusions during dexterous manipulation; (3) The unconstrained interaction causes more variations in hand pose and contact than in existing datasets [8, 22, 23, 41] (see Fig. 2).

As an initial step towards addressing these challenges, and to provide baselines for future work, we introduce ArcticNet to reconstruct the motions of two hands and an articulated object from a video. ArcticNet uses an encoder-decoder architecture to estimate parameters of the MANO hand model [53] for the two hands, and our articulated object model. We experiment with two variations of ArcticNet: a single-frame model and a temporal model with a recurrent architecture inspired by [34]. We provide qualitative and quantitative results for future comparison.

When studying hand-object interaction, contact is important [21, 77]. Some approaches [26, 77] explore the task of binary contact estimation from a single RGB image. In the two-handed manipulation setting, hands can be near the object but not in contact. To understand the dynamic, relative spatial configuration between hands and objects in more detail, even when not in contact, we propose the general task of *interaction field estimation* from RGB images. The goal is to estimate, for each hand vertex, the shortest distance to the object mesh and vice versa (see Fig. 6 for a visualization). We introduce a baseline, InterField, for this task and benchmark both a single-frame and a recurrent version of InterField on ARCTIC for future comparison.

In summary, our contributions are as follows: (1) We present ARCTIC, the first large-scale dataset of two hands that *dexterously* manipulate *articulated* objects, with multi-view RGB images paired with accurate 3D meshes; (2) We introduce two novel tasks of consistent motion reconstruction and interaction field estimation to study the physically consistent motion of hands and articulated objects; (3) We provide baselines for both tasks on ARCTIC.

2. Related Work

Human-object datasets: Several datasets [1, 7, 46, 61, 70, 73] contain images of human-object interaction, but here we focus on large-scale data [3, 19, 22, 25, 27, 55, 89] that facilitates machine learning. There are three categories. *(1) Human body with rigid objects:* Bhatnagar *et al*. [3] and Huang *et al*. [27] introduce image datasets for human body interaction with big objects. Compared to ours, [3] do not capture the hands. Huang *et al*. [27] capture hands and body using a multi-view RGB-D setup while ours is captured using a MoCap setup for more accurate 3D data. Compared to both, we have dexterous bimanual manipulation, dynamic hand-object contact, and articulated objects. GRAB [65] contains detailed human-object interaction but no images, while BEDLAM [4] contains videos with ground-truth humans but no object interaction. *(2) Single hand with rigid objects:* Most hand-object datasets [6, 8, 19, 22, 25, 41] consist of single-hand grasping interaction. However, hand poses in grasping interaction are mostly static, with relatively little pose variation over time. Hampali *et al*. [22] use a multi-RGB-D system and fit both MANO and YCB object meshes with sequence-level fitting and contact constraints. *(3) Two hands with rigid objects:* Kwon *et al*. [37] and Hampali *et al*. [23] present two-hand datasets interacting with rigid objects. Compared to (2) and (3), our dataset has 3D annotations of the full human body, both hands, and articulated objects. We go beyond grasping and focus on less constrained dexterous bimanual manipulation. We discuss the comparison between ours (ARCTIC) and existing hand-object datasets [8, 22, 23, 37, 41] in Sec. 3.1.

Estimating 3D hands and objects from RGB images: Monocular RGB 3D hand reconstruction has a long history since Rehg and Kanade [51]. Most work in the literature focuses on hand-only reconstructions [5, 17, 25, 28, 38, 44, 45, 57–60, 71, 80, 83, 87, 87, 88]. Zimmermann *et al*. [88] use a deep convolutional network for 3D hand pose estimation via a multi-stage approach. Spurr *et al*. [59] introduce biomechanical constraints to regularize hand pose prediction. Ziani *et al*. [87] use a self-supervised time-contrastive formulation to improve smoothness for hand motion reconstruction. Recently, there has been increased interest in hand-object reconstruction from RGB images [12, 21, 24, 25, 40, 66, 77, 86]. Tekin *et al*. [66] infer 3D control points for both the hand and the object in videos, using a temporal model to propagate information across time. Hasson *et al*. [25] render synthetic images and train a neural network to regress a static grasp of a 3D hand and a rigid object, using full supervision together with contact losses. Corona *et al*. [12] estimate MANO grasps for objects from an image, by first inferring the object shape and a rough hand pose, which is refined via contact constraints and an adversarial prior. Liu *et al*. [40] use a transformer-based contextual-reasoning module that encodes the synergy between hand

dataset	real images	# number of: img	# number of: view	ego-centric	image resol.	articulated objects	both hands	human body	dexterous manipulation	annot. type
FreiHand [89]	✓	37k	8	✗	224×224	✗	✗	✗	✗	semi-auto
ObMan [25]	✗	154k	1	✗	256×256	✗	✗	✗	✗	synthetic
FHPA [19]	✓	105k	1	✓	1920×1080	✗	✗	✗	✗	magnetic
HO3D [22]	✓	78k	1-5	✗	640×480	✗	✗	✗	✗	multi-kinect
ContactPose [6]	✓	2.9M	3	✗	960×540	✗	✗	✗	✗	multi-kinect
GRAB [65]	-	-	-	-	-	✗	✓	✓	✗	mocap
DexYCB [8]	✓	582k	8	✗	640×480	✗	✗	✗	✗	multi-manual
H2O [37]	✓	571k	5	✓	1280×720	✗	✓	✗	✗	multi-kinect
H2O-3D [23]	✓	76k	5	✗	640×480	✗	✓	✗	✗	multi-kinect
HOI4D [41]	✓	2.4M	1	✓	1280×800	✓	✗	✗	✗	single-manual
ARCTIC (Ours)	✓	2.1M	9	✓	2800×2000	✓	✓	✓	✓	mocap

Table 1. **Comparison of our ARCTIC dataset with existing datasets**. The keyword "single/multi-manual" denotes whether single or multiple views being used to annotate manually.

and object features, and has higher responses at contact regions. Zhou *et al*. [84] learn an interaction motion prior to denoise motion predicted from an off-the-shelf single-frame hand-object reconstruction method. None of these methods deal with articulated objects, which result in complex hand-object interactions.

Human-object contact detection: Contact has been shown important for: pose taxonomies [2, 18, 29], pose estimation [21, 22, 25, 61, 69, 73, 77], in-hand scanning [72, 82], and grasp synthesis [21, 31, 65, 77]. Many methods [21, 22, 61, 69, 73] use the proximity between the 3D hand/body and object meshes to estimate contacts and regularize pose estimation based on these. Three main categories for contact estimation exist: 1) directly from meshes; 2) on the image pixel space from RGB images; 3) binary contact in 3D space from RGB images. Grady *et al*. [21] take off-the-shelf regressors to estimate grasping hand and object meshes, use these meshes to predict contacts on the objects provided by [6], and leverage contacts to refine the grasp. Their recent dataset [20] contains both contact and pressure between a hand and a flat sensor surface. Tripathi *et al*. [68] infer pressure from body-scene contact. Narasimhaswamy *et al*. [47] and Shan *et al*. [56] infer bounding boxes for hands in contact on the input RGB image. Chen *et al*. [9] infer human-scene contact on pixels. Rogez *et al*. [52] learn to infer contacts from the image using synthetic data, while Pham *et al*. [49] use real contact data captured with instrumented objects. Unlike others, [52] and [49] estimate 3D binary contact from RGB images but the former does not generalize well to real images and the latter uses a classical approach due to the limited amount of data. BSTRO estimates contact on the 3D body from an image but does not estimate 3D hand or object pose [26]. Hi4D [78] provides ground-truth contact for close human interaction. In contrast, our task of interaction field estimation goes beyond binary contact to model the dense relative distances between hands and objects. Thanks to our dexterous manipulation, ARCTIC contains fast changing hand-object contact.

3. ARCTIC Dataset

Overview: To allow the study of object articulation with hands in motion, we construct ARCTIC, a video dataset with accurate 3D annotation for hands and articulated objects. ARCTIC contains 339 sequences of dexterous manipulation of 11 articulated objects by 10 subjects (5 fe-/males). The dataset consists of 2.1M RGB images from 8 static views and 1 egocentric view, paired with 3D hand and object meshes. To capture different interaction modes, we ask our subjects to either "use" (1.7M images) or "grasp" (457K images) the objects. Depth images of the two hands, the human body, and objects can be rendered from ARCTIC (see SupMat).

3.1. Data Characteristics

Dataset features comparison: Table 1 compares ARCTIC with existing hand-object datasets. ARCTIC is the only dataset that contains both hands, the full human body (in SMPL-X [48]) and articulated objects. ARCTIC provides calibrated cameras (8 allocentric and 1 egocentric) with high-resolution images, enabling the study of monocular, multi-view and egocentric reconstruction settings. Importantly, ARCTIC is a motion dataset that focuses on bi-manual dexterous manipulation, meaning that subjects can freely interact with objects using both hands. In contrast, existing hand-object datasets focus single-hand grasping [8, 22, 25] and the movement is often controlled [23, 37]. GRAB [65] has fast motion by using a similar MoCap setup but captures only rigid objects and does not have images. HOI4D [41] is the only hand-object dataset that contains articulated objects, but it contains only a single view, does not capture the full human body, has a single hand, and mainly focuses on grasping. Crucially, their hand data is captured from only a single egocentric view, which introduces ambiguity for the occluded fingers.

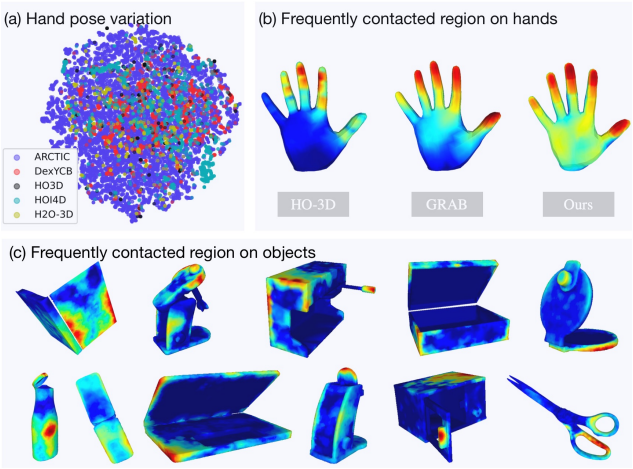


Figure 2. **Hand pose and contact variations in datasets**. (a) T-SNE clustering of hand poses in different datasets. The plot shows that ARCTIC has a significantly larger range of poses than all existing datasets. (b) Frequently contacted regions for hands in HO-3D [22], GRAB [65], and ARCTIC. As seen with the broader heatmap spread on the hands, ARCTIC has higher contact diversity. (c) Frequently contacted areas on our objects.

Capture setup comparison: Capturing dexterous manipulation while maintaining the quality of 3D annotation is extremely challenging due to fast motion and heavy occlusion during the interaction. In particular, the joints of a hand often have significant self-occlusion. The occlusion is even more severe when a hand interacts with objects and when there are multiple hands [44]. Existing hand-object datasets [8,22,23,37,41] are captured with $1-8$ commodity RGB-D cameras, which is insufficient to eliminate occlusion. As a result, their hand-object motion is often slow and they mainly focus on grasping interaction. To reduce occlusion and to enable the capture of dexterous manipulation, we construct our dataset using an accurate Vicon MoCap setup with 54 high-end infrared Vantage-16 cameras [75]. To show our dexterous motion, and to compare 3D annotation quality between datasets, see our project page video.

Hand pose and contact variations: Figure 2a compares different hand-object datasets [8,22,23,41] in terms of hand pose variations by showing a T-SNE clustering [74] of 3D hand joints. The plot reveals that our dataset (shown in blue) has a significantly larger hand pose diversity than others. This is due to the unconstrained nature of ARCTIC in which the subjects dexterously and dynamically *manipulate* the object (see project page video). The figure also shows frequently in-contact regions on hands (b) and objects (c) in the ARCTIC dataset. We generate the contact heatmaps following GRAB's [65] approach, by integrating per-frame binary contact labels for vertices over all sequences. "Hotter" regions denote a higher chance of being in contact while "cooler" regions denote lower chance of contact. Similar to HO-3D [22] and GRAB [65], finger tips in our dataset are most likely to be in contact with objects. However, thanks to the dexterous manipulation it contains, ARCTIC has higher contact likelihood in the palm region than other datasets, hence the heatmaps appear more "spread out". For regular-sized everyday objects, such as the ketchup bottle, the contact regions "agree" with our usual interaction with them. For smaller toy objects like the waffle iron, subjects are likely to pick up the object and support it with one hand, leading to "hot" regions on the bottom of the object.

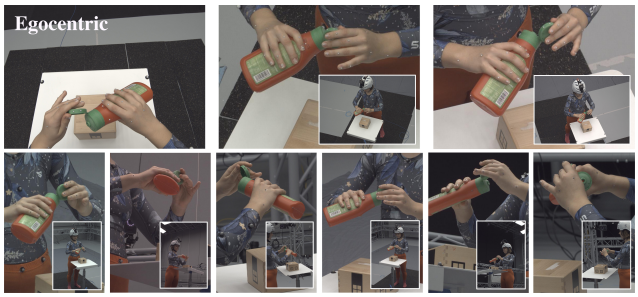


Figure 3. **Our camera views**. We capture high resolution images in 8 static allocentric and 1 moving egocentric views. Here we show zoomed-in crops and the original images.

3.2. Acquisition Setup

We detail our motion capture (MoCap) setup to acquire 3D surfaces of strongly interacting hands and articulated objects. We synchronize a MoCap system with a multi-view RGB system. See SupMat for the marker sets. With the latter we capture RGB videos from 8 static allocentric views and 1 moving egocentric view at 30 FPS (see Fig. 3). The capture pipeline has five steps: (1) obtaining the 3D template geometry of the subjects and objects, (2) estimating the rotation axis for articulated objects, shown in SupMat, (3) capturing interaction using marker-based MoCap together with calibrated and synchronized video, (4) solving for the poses of the body, hands, and objects from MoCap markers following [42, 65], and (5) computing hand-object contact based on proximity, shown in SupMat.

Obtaining canonical geometry: We obtain the ground-truth (GT) hand and body shape of each subject in a canonical T-Pose using 3D scans from a 3dMD [67] scanner. We register SMPL-X [48] to 3D scans at different time steps in varying poses and construct a personalized 3D template of each subject. See the SupMat for details of the template creation. To obtain object geometries, we scan each object using an Artec 3D hand-held scanner in a pre-defined pose. We separate each scanned object mesh into two articulated parts in Blender. See SupMat for all 11 articulated objects.

Capturing human-object interaction: To ensure accuracy, we perform full-body, hand and object tracking using a Vicon MoCap system with 54 infrared Vantage-16 cameras [75] to minimize the issues with occlusion. To capture

usable RGB images alongside the MoCap data, we balance the trade-off between accuracy and marker intrusiveness by using small hemispherical markers with 1.5mm radius on the hands and objects. The markers are placed on the dorsal side of the hand to not encumber participants during natural hand-object interaction, similar to GRAB [65]. While our focus is on hands, we retrieve full-body pose estimates as they provide more reliable global rotations and translations for each hand. Therefore, we fit SMPL-X [48] to the observed markers to attain realistic wrist articulations, as MANO contains no wrist articulation.

Obtaining surfaces from MoCap: Following [42, 65], we associate MoCap marker positions with their corresponding subject/object vertices in the geometries obtained in canonical spaces. We first pick initial guesses of marker-to-vertex correspondence on the subject/object meshes and use MoSh++ [42] to refine the correspondence. To obtain the full-body and hand surface that explain the MoCap data, we optimize SMPL-X pose using each subject's SMPL-X template to minimize the distance between the markers and their correspondences on the SMPL-X mesh.

The articulated object surface is parameterized by the 6D pose of each object's base part and an 1D articulation relative to a canonical pose. We obtain the 6D pose of the object base for each MoCap frame by solving for the rigid transformation between the MoCap markers of the object base at that frame, and the object vertices corresponding to the markers in the object canonical space. The 1D articulation is computed according to the estimated rotation axis (see SupMat) and a pre-defined rest pose.

4. Evaluation Protocol

Data split: We split the data by subjects, 8 subjects for training, 1 for validation (male) and 1 for testing (female). To ensure gender balance in evaluation, we use one male and one female subject. With this same split, we establish two protocols: an allocentric protocol (**allo**) and an egocentric protocol (**ego**). The former protocol lets us study our tasks in the 3rd-person, while the latter is similar to 1st-person views in a mixed-reality setting. In the allocentric protocol, during training and evaluation, the model only has access to images from the allocentric views. In the egocentric protocol, to provide additional training images, we allow models access to images from all views of the training split, but in evaluation, only egocentric images are used. Further information can be found in SupMat.

Metrics for consistent motion reconstruction: Our goal is to reconstruct the 3D motion of the hands and an articulated object during dexterous manipulation from a video. Importantly, our focus extends beyond hand-object poses and we require the reconstructed meshes to have accurate hand-object contact (CDev), and smooth motion (ACC). Further, when a hand moves or articulates an object, vertices of the hand and the object in stable contact should move together (MDev). To this end, we define the following metrics:

- **Contact Deviation (CDev):** For a frame, suppose $\{(\mathbf{h}_i, \mathbf{o}_i)\}_{i=1}^{C}$ are C pairs of in-contact hand-object vertices ($< 3mm$ distance in ground-truth), and $\{(\hat{\mathbf{h}}_i, \hat{\mathbf{o}}_i)\}_{i=1}^{C}$ are the corresponding predictions. CDev is defined as the average distance between $\hat{\mathbf{h}}_i$ and $\hat{\mathbf{o}}_i$ in millimeters:
$$\frac{1}{C}\sum_{i=1}^{C} ||\hat{\mathbf{h}}_i - \hat{\mathbf{o}}_i|| \quad (1)$$
This metric reflects how much the hand vertices deviate from the supposed contact vertices on the object in the prediction.
- **Motion Deviation (MDev):** Given a ground-truth sequence of a hand and an object, we denote vertex i of the hand and vertex j of the object at frame t as $\mathbf{h}_i^t$, $\mathbf{o}_j^t$ respectively. We use (i, j, m, n) to denote $\mathbf{h}_i^t$ has stable contact with $\mathbf{o}_j^t$ during a window from frame m to frame n, and they do not have contact at time $m-1$ and $n+1$ (*i.e.*, longest contact window). Hand-object vertex indices (i, j) have stable contact in a window (m, n) if they are close within a threshold α for every frame in the window:
$$\forall t \in \{m, \cdots, n\}, \left\|\mathbf{h}_i^t - \mathbf{o}_j^t\right\| \leq \alpha\,. \quad (2)$$
Given the above definition, we extract a set of tuples $\{(i, j, m, n)\}$ from each GT sequence. When two hand-object vertices $\mathbf{h}_i^t$, $\mathbf{o}_j^t$ are in stable contact within a window, they should move in the same direction in consecutive frames. To measure this, we define the motion deviation for a tuple (i, j, m, n) of the predicted hand-object sequence $\hat{\mathbf{h}}$ and $\hat{\mathbf{o}}$ as
$$\frac{1}{n-m}\sum_{t=m+1}^{n} ||\delta\hat{\mathbf{h}}_i^t - \delta\hat{\mathbf{o}}_j^t|| \quad (3)$$
where $\delta\hat{\mathbf{h}}_i^t = \hat{\mathbf{h}}_i^t - \hat{\mathbf{h}}_i^{t-1}$ and $\delta\hat{\mathbf{o}}_j^t = \hat{\mathbf{o}}_j^t - \hat{\mathbf{o}}_j^{t-1}$. Intuitively, this measures the disagreement in the moving direction between consecutive frames of in-contact hand-object vertices in the window (m, n). We only consider longer motions by using windows with at least 0.5 second or 15 frames (*i.e.*, $n - m + 1 \geq 15$) and we choose $\alpha = 3mm$ to detect a sufficient number of windows. We compute this metric for all detected windows and average over them.
- **Acceleration Error (ACC):** Following [34], we report acceleration error in m/s^2 to measure the smoothness of the reconstruction, calculated as the difference in acceleration between the ground-truth and predicted vertex sequences for each hand and the object. We subtract the root for each entity before computing the acceleration [34]. The root for the object is defined as the center of an object's base. Note that we report this error in m/s^2, while [34] reports mm/s^2. See SupMat for more details.

Apart from motion and contact, we need metrics to measure hand and object poses, and their relative translations:

- **Mean Per-Joint Position Error (MPJPE):** the L2 distance (mm) between the 21 predicted and ground-truth joints for each hand after subtracting its root.
- **Average Articulation Error (AAE):** the average absolute error between the predicted degree of articulation and the ground-truth.
- **Success Rate:** Following [62, 79], to measure object reconstruction quality, we use a success rate metric that is independent of the object size. It is the percentage of predicted object vertices having L2 error to the ground-truth that is less than 5% of the object diameter:

$$\frac{1}{V_o}\sum_{i=1}^{V_o} \mathbb{1}(\|\mathbf{o}_i - \hat{\mathbf{o}}_i\| < 0.05D) \times 100\% \quad (4)$$

  where D, V_o, $\mathbf{o}_i$, $\hat{\mathbf{o}}_i$ are the diameter, the number of object vertices, ground-truth and predicted object vertices, and $\mathbb{1}(\cdot)$ is the indicator function. To decouple the effect of root estimation, we subtract the predicted and the ground-truth vertices by their object roots respectively. The root is the center of each object's base.
- **Mean Relative-Root Position Error (MRRPE):** Following [17, 44], to measure the root translation of between hand-hand and hand-object, we use this metric to measure the relative root translation between two entities a and b in the scene,

$$\text{MRRPE}_{a\to b} = \left\| (\mathbf{J}_0^a - \mathbf{J}_0^b) - \left(\hat{\mathbf{J}}_0^a - \hat{\mathbf{J}}_0^b\right) \right\|_2, \quad (5)$$

  where $a \in \{l, r, o\}$ and $b \in \{l, r, o\}$ and l, r, o denote the left hand, right hand, and the object, $\mathbf{J}_0 \in \mathbb{R}^3$ is the ground-truth root joint location and $\hat{\mathbf{J}}_0$ the predicted one. A graphical illustration of this metric can be found in SupMat.

Metrics for interaction field estimation: In this task, given images from a video, for each hand vertex i, we estimate its shortest distance $\hat{\mathbf{F}}_i^{r\to o} \in \mathbb{R}$ to the object (*i.e.*, the distance field from a hand to the object) and vice versa. Taking the field from the right hand to the object as an example, to quantify, we measure the average error between the predicted distances $\hat{\mathbf{F}}_i^{r\to o}$ and the ground-truth distances $\mathbf{F}_i^{r\to o}$ in millimeters, which we call average distance error. The error is computed as:

$$\frac{1}{V_r}\sum_{i=1}^{V_r} |\mathbf{F}_i^{r\to o} - \hat{\mathbf{F}}_i^{r\to o}| \quad (6)$$

where V_r is the number of right-hand vertices. To measure smoothness, we estimate the distance field for every frame in each sequence. We then compute the acceleration sequence for the predicted field sequence. The acceleration error is computed as the average absolute difference between predicted and ground-truth acceleration sequences. See SupMat for the formula of acceleration error.

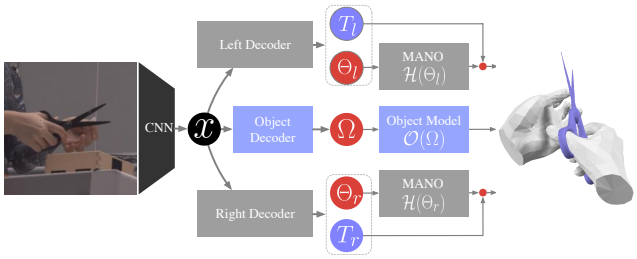


Figure 4. **ArcticNet-SF architecture**. The CNN encoder yields image features x. The hand decoders predict MANO parameters $\mathbf{\Theta}_l, \mathbf{\Theta}_r$ and their translation $\mathbf{T}_l, \mathbf{T}_r$ while the object decoder estimates the articulated object pose $\mathbf{\Omega}$ consisting of the articulation, rotation and translation. With parametric models of hands $\mathcal{H}(\mathbf{\Theta})$ and articulated objects $\mathcal{O}(\mathbf{\Omega})$, we obtain 3D meshes for the two hands and the articulated object .

5. Baselines and Experiments

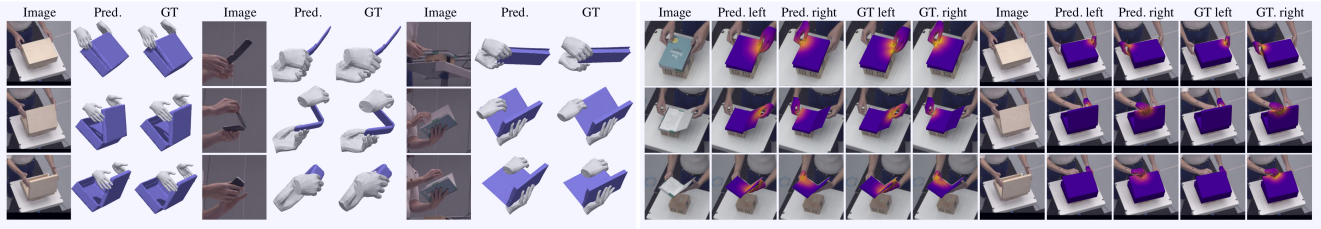
We present two tasks on ARCTIC: consistent motion reconstruction and interaction field estimation. For consistent motion reconstruction, we reconstruct the 3D motion of two hands and an articulated object from a video. For interaction field estimation, given a video, we estimate, for each hand vertex, the closest distance to the object and vice versa. Here we detail and evaluate our baselines in the two tasks to lay the foundation for future comparison.

5.1. Consistent motion reconstruction

Problem formulation: Given a video, our goal is to reconstruct the 3D motion of a subject's two hands and an articulated object in dexterous manipulation for every frame. Our emphasis is to require the reconstructed hand-object meshes to be in temporally-consistent hand-object contact and motion during object articulation and manipulation.

Parametric models: For brevity, we use l, r, and o to denote the left hand, the right hand and the object. For hands, we use MANO [53] to represent the hand pose and shape by $\mathbf{\Theta} = \{\theta, \beta\}$, which consists of parameters for the pose $\theta \in \mathbb{R}^{48}$ (with global orientation) and the shape $\beta \in \mathbb{R}^{10}$. The MANO model maps $\mathbf{\Theta}$ to a shaped and posed 3D mesh $\mathcal{H}(\theta, \beta) \in \mathbb{R}^{778\times 3}$. The 3D joint locations $\mathbf{J} = W\mathcal{H} \in \mathbb{R}^{\mathrm{J}\times 3}$ are obtained using a pre-trained linear regressor W. For each object, we construct a 3D model $\mathcal{O}(\cdot)$ using the scanned object mesh, the estimated rotation axis, and the marker-vertex correspondences estimated in Sec. 3.2. The function takes as inputs the articulated object pose, $\mathbf{\Omega}$, and outputs a posed 3D mesh, $\mathcal{O}(\mathbf{\Omega}) \in \mathbb{R}^{\mathrm{V}\times 3}$, where V denotes the object's number of vertices. The object pose, $\mathbf{\Omega} \in \mathbb{R}^7$, consists of the 1D rotation (radians) for articulation, $\omega \in \mathbb{R}$, and the 6D object rigid pose, *i.e.*, its rotation, $\mathbf{R}_o \in \mathbb{R}^3$, and translation, $\mathbf{T}_o \in \mathbb{R}^3$.

Baselines: We introduce ArcticNet to estimate the poses of the two hands and the articulated object from RGB images.



(a) Consistent motion reconstruction (b) Interaction field estimation

Figure 5. **Qualitative results of ArcticNet-LSTM (a) and InterField-LSTM (b)**. Best viewed in color and zoomed in. See SupMat for results of ArcticNet-SF and InterField-SF.

		Contact and Relative Position		Motion		Hand	Object	
Splits	Method	CDev$_{ho}$ $[mm]\downarrow$	MRRPE$_{rl/ro}$ $[mm]\downarrow$	MDev$_{ho}$ $[mm]\downarrow$	ACC$_{h/o}$ $[m/s^2]\downarrow$	MPJPE$_{h}$ $[mm]\downarrow$	AAE $[°]\downarrow$	Success Rate $[\%]\uparrow$
Allo. Val	ArcticNet-SF	41.4	50.1/37.6	10.4	6.6/8.8	23.0	5.9	71.8
	ArcticNet-LSTM	**38.8**	**47.1/36.8**	**8.9**	**5.6/6.9**	**22.9**	**5.8**	**74.9**
Allo. Test	ArcticNet-SF	41.6	52.4/**37.5**	10.4	5.7/7.6	**21.5**	5.4	71.4
	ArcticNet-LSTM	**38.9**	**49.2**/37.7	**9.3**	**5.0/6.1**	**21.5**	**5.2**	**73.5**
Ego. Val	ArcticNet-SF	**44.1**	**33.9/36.8**	11.8	6.3/11.3	**22.9**	**8.0**	59.0
	ArcticNet-LSTM	44.5	39.3/39.0	**8.1**	**4.3/7.2**	23.8	**8.0**	**59.1**
Ego. Test	ArcticNet-SF	44.7	**28.3**/36.2	11.8	5.0/9.1	**19.2**	**6.4**	**53.9**
	ArcticNet-LSTM	**43.3**	31.8/**35.0**	**8.6**	**3.5/5.7**	20.0	6.6	53.5

Table 2. **Comparison of two reconstruction baselines**. Contact and relative position metrics measure hand-object contact and relative root position prediction. Motion metrics reflect motions with temporally-consistent contact and smoothness. Hand and object metrics show root-relative reconstruction error. See Sec. 4 for metric details. We use l, r, o to denote the left, the right hand, and the object. To simplify the results, we average left and right hand metrics into one hand (denoted by h). For example, CDev$_{ho}$ is the contact deviation between a hand and the object averaged over the two hands; MRRPE$_{rl/ro}$ denotes MRRPE$_{r\rightarrow l}$ and MRRPE$_{r\rightarrow o}$ between the slash.

We benchmark two versions of ArcticNet: a single-frame model (ArcticNet-SF), and a model with a recurrent architecture (ArcticNet-LSTM). The LSTM baseline is used to allow a joint reasoning of hand and articulated object motions. Figure 4 summarizes the architecture of ArcticNet-SF. Inspired by Hasson *et al*. [24, 25], we use an encoder-decoder architecture. In particular, the CNN encoder takes in the input image and produces image features $\mathbf{x}$. The features are used by the hand decoders to estimate the parameters for the left and right hands, $\mathbf{\Theta}_l$ and $\mathbf{\Theta}_r$, as well as the translations for the two hands, $\mathbf{T}_l$ and $\mathbf{T}_r$. Similarly, the object decoder predicts the articulated object pose, $\mathbf{\Omega}$. We use axis-angle for rotation and use the weak perspective camera model to estimate the translations [5, 30, 35, 54, 83]. The ArcticNet-LSTM model has the same architecture as ArcticNet-SF, except that it has an LSTM network to aggregate image features from multiple frames before passing them to the regression heads. We train the models with ground-truth 3D keypoints, 2D projected keypoints, and the parameters of the hand and the object models. We show details of the model and the training procedure in SupMat.

Results: Figure 5a visualizes the predictions of one of our baselines, ArcticNet-LSTM with [32]. To see qualitative results of ArcticNet-SF, we refer to the SupMat. Table 2 shows the quantitative evaluation of the two baseline models on ARCTIC. The results show that, overall, the ArcticNet-LSTM model has temporally more consistent contact (CDev), and motion (MDev) between the hands and objects. Further, it has smoother motion (ACC). This demonstrates that temporal modelling is important for spatio-temporally consistent hand-object motion and contact. See Sec. 4 for metric details.

5.2. Interaction field estimation

Existing contact detection methods mainly focus on binary contact estimation [21, 77]. In two-handed dexterous interactions, hands can be near the object, but not always in contact. We define a general task of interaction field estimation to capture the relative spatial relations between hands and the object even when not in contact.

Problem formulation: We define an interaction field $\boldsymbol{F}^{a\rightarrow b} \in \mathbb{R}^{V_a}$ as the distance to the closest vertex on the mesh $\mathbf{M}_b$ for all vertices in mesh $\mathbf{M}_a$ where V_a (or V_b) is the number of vertices in mesh $\mathbf{M}_a$ (or $\mathbf{M}_b$). Formally,

$$\boldsymbol{F}_i^{a\rightarrow b} = \min_{1\leqslant j\leqslant V_b} ||\mathbf{v}_i^a - \mathbf{v}_j^b||_2, \quad 1 \leqslant i \leqslant V_a \qquad (7)$$

where $\mathbf{v}_k^m \in \mathbb{R}^3$ represents the k-th vertex of mesh $\mathbf{M}_m$. We define our task to estimate the interaction fields $\boldsymbol{F}^{l\rightarrow o}$, $\boldsymbol{F}^{r\rightarrow o}$, $\boldsymbol{F}^{o\rightarrow l}$, and $\boldsymbol{F}^{o\rightarrow r}$ for each image. In other words, for each vertex of each hand we aim to infer the closest distance to the object and vice-versa.

Splits	Method	Average Distance Error $[mm]\downarrow$	ACC $[m/s^2]\downarrow$
Allo. Val	InterField-SF	9.6/9.9	3.0/2.9
	InterField-LSTM	**9.0/8.9**	**2.1/2.0**
Allo. Test	InterField-SF	9.0/10.0	2.7/2.7
	InterField-LSTM	**8.7/9.1**	**1.9/1.9**
Ego. Val	InterField-SF	8.8/9.2	2.4/2.3
	InterField-LSTM	**8.4/8.9**	**2.1/2.0**
Ego. Test	InterField-SF	8.2/9.2	2.1/2.0
	InterField-LSTM	**8.0/9.1**	**1.8/1.8**

Table 3. **Comparison of two field estimation baselines**. To simplify the evaluation, we average metrics for the two hands into one. The slashes denote the average distance error and the acceleration error for hand-to-object/object-to-hand.



Figure 6. **InterField-SF architecture**. We concatenate image features $\mathbf{x}$ to each subsampled hand-object vertex in canonical pose. The concatenated vectors are passed through a PointNet and then regressed to distance values. The interaction field is visualized as a heatmap for each entity (bright: closest vertex is near).

Baselines: We present InterField to estimate the interaction field from RGB images. We benchmark two versions of InterField: a single-frame (InterField-SF) and a temporal baseline (InterField-LSTM). The temporal model lets us evaluate the benefits of temporal information. Figure 6 outlines the framework of InterField-SF. Suppose that we estimate the field $\hat{\boldsymbol{F}}^{l\to o}$. We first extract image features $\mathbf{x} \in \mathbb{R}^{\mathrm{d}}$ via a CNN backbone. Next, we concatenate $\mathbf{x}$ to each sub-sampled vertex of the left hand (l) in its canonical pose to obtain $\mathbf{p}_i = [\mathbf{x}; \mathbf{v}_i] \in \mathbb{R}^{\mathrm{d}+3}$ for all $1 \leqslant i \leqslant \bar{V}_l$ where $\bar{V}_l$ denotes the number of subsampled vertices. All points $\mathbf{p}_i$ are fed to a PointNet [50] followed by a regression head that estimates the distance. The predicted distances are upsampled to the full mesh. For efficiency, we use subsampled vertices for the PointNet and upsample for regression. The remaining interaction fields are estimated via the same network with a shared CNN and PointNet but different heads. InterField-LSTM follows the same formulation except it has an LSTM to aggregate image features in a temporal window to jointly reason about hand-object motion. See more training and baseline details in SupMat.

Results: Figure 5b shows qualitative samples of InterField-LSTM. The predicted values are visualized as heatmaps over the meshes of the respective hands or objects. A "hotter" region denotes closer distances. Note that the ground-truth meshes are only used for visualization; they are not network inputs. We find that the predicted fields correlate well with the ground truth. Table 3 shows the performance of our baselines. The results show that modeling the hand-object interaction field over time yields more accurate results (see distance error), and smoother predictions (ACC).

6. Conclusions

We introduce ARCTIC, the first dataset with two hands dexterously manipulating articulated objects that includes high-quality 3D ground-truth for hands, and objects together with synchronized video. ARCTIC has a total of 2.1M RGB images from 8 static views and 1 egocentric view of 10 subjects interacting with 11 articulated objects. We present two tasks on ARCTIC. First is *consistent motion reconstruction*. Given a video, we reconstruct two hands and an articulated object in 3D for every frame, such that their motions are spatio-temporally consistent. The second task is *interaction field estimation*, where we estimate dense relative hand-object distances from images in a video. We present two baselines ArcticNet and InterField for the two tasks respectively, and evaluate them on ARCTIC to lay the foundation for future work.

Future directions: ARCTIC can enable a range of tasks related to hand manipulation with object articulation. First, methods for generating hand-object interaction focus on generating grasps of rigid objects [11, 31], but less work has been done on generating dexterous bimanual manipulation motion with objects [10, 81] and prior work does not generate interaction with articulated objects (*e.g.*, "cutting with scissors"). ARCTIC can enable these new generation tasks, and extend them to the full-body [64] with our SMPL-X ground-truth. Second, we introduce tasks of consistent motion reconstruction and interaction field estimation. Future work could leverage the interaction field representation for pose estimation to improve hand-object contact in reconstruction. Finally, articulated object pose estimators [39] from depth images do not consider humans in the scene. The rendered depth images in ARCTIC can be used to benchmark such methods in more realistic settings.

Acknowledgements: The authors deeply thank Tsvetelina Alexiadis (TA) for trial coordination; Markus Höschle (MH), Senya Polikovsky, Matvey Safroshkin, Tobias Bauch (TB) for the capture setup; MH, TA, Galina Henz for data capture; Giorgio Becherini, Nima Ghorbani for MoSh++; Priyanka Patel for alignment; Leyre Sánchez Vinuela, Andres Camilo Mendoza Patino, Mustafa Alperen Ekinci for data cleaning; TB for Vicon support; MH, Jakob Reinhardt for object scanning; Taylor McConnell for Vicon support, and data cleaning coordination; Benjamin Pellkofer for IT support. We also thank Xu Chen, Adrian Spurr, Jie Song for insightful discussion. OT and DT were partially funded by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039B. DT's work was partially performed at the MPI-IS.

Disclosure: https://files.is.tue.mpg.de/black/CoI_CVPR_2023.txt

References

[1] Luca Ballan, Aparna Taneja, Jürgen Gall, Luc Van Gool, and Marc Pollefeys. Motion capture of hands in action using discriminative salient points. In *European Conference on Computer Vision (ECCV)*, pages 640–653, 2012. 2

[2] Keni Bernardin, Koichi Ogawara, Katsushi Ikeuchi, and Ruediger Dillmann. A sensor fusion approach for recognizing continuous human grasping sequences using hidden markov models. *Transactions on Robotics*, 21(1):47–57, 2005. 3

[3] Bharat Lal Bhatnagar, Xianghui Xie, Ilya A Petrov, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. BEHAVE: Dataset and method for tracking human object interactions. In *Computer Vision and Pattern Recognition (CVPR)*, pages 15935–15946, 2022. 2

[4] Michael J. Black, Priyanka Patel, Joachim Tesch, and Jinlong Yang. BEDLAM: A dataset of bodies exhibiting detailed lifelike animated motion. In *Computer Vision and Pattern Recognition (CVPR)*, June 2023. 2

[5] Adnane Boukhayma, Rodrigo de Bem, and Philip H. S. Torr. 3D hand shape and pose from images in the wild. In *Computer Vision and Pattern Recognition (CVPR)*, pages 10843–10852, 2019. 2, 7, 16

[6] Samarth Brahmbhatt, Chengcheng Tang, Christopher D. Twigg, Charles C. Kemp, and James Hays. ContactPose: A dataset of grasps with object contact and hand pose. In *European Conference on Computer Vision (ECCV)*, volume 12358, pages 361–378, 2020. 2, 3

[7] Zhe Cao, Ilija Radosavovic, Angjoo Kanazawa, and Jitendra Malik. Reconstructing hand-object interactions in the wild. In *International Conference on Computer Vision (ICCV)*, pages 12417–12426, 2021. 2

[8] Yu-Wei Chao, Wei Yang, Yu Xiang, Pavlo Molchanov, Ankur Handa, Jonathan Tremblay, Yashraj S. Narang, Karl Van Wyk, Umar Iqbal, Stan Birchfield, Jan Kautz, and Dieter Fox. DexYCB: A benchmark for capturing hand grasping of objects. In *Computer Vision and Pattern Recognition (CVPR)*, pages 9044–9053, 2021. 1, 2, 3, 4

[9] Yixin Chen, Sai Kumar Dwivedi, Michael J. Black, and Dimitrios Tzionas. Detecting human-object contact in images. In *Computer Vision and Pattern Recognition (CVPR)*, June 2023. 3

[10] Yuanpei Chen, Yaodong Yang, Tianhao Wu, Shengjie Wang, Xidong Feng, Jiechuang Jiang, Stephen Marcus McAleer, Hao Dong, Zongqing Lu, and Song-Chun Zhu. Towards human-level bimanual dexterous manipulation with reinforcement learning. *arXiv preprint arXiv:2206.08686*, 2022. 8

[11] 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)*, pages 20545–20554, 2022. 8

[12] Enric Corona, Albert Pumarola, Guillem Alenyà, Francesc Moreno-Noguer, and Grégory Rogez. GanHand: Predicting human grasp affordances in multi-object scenes. In *Computer Vision and Pattern Recognition (CVPR)*, pages 5030–5040, 2020. 2

[13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition (CVPR)*, pages 248–255. Ieee, 2009. 15

[14] Bardia Doosti, Shujon Naha, Majid Mirbagheri, and David J. Crandall. HOPE-Net: A graph-based model for hand-object pose estimation. In *Computer Vision and Pattern Recognition (CVPR)*, pages 6607–6616, 2020. 19

[15] Yuval Eldar, Michael Lindenbaum, Moshe Porat, and Yehoshua Y. Zeevi. The farthest point strategy for progressive image sampling. In *International Conference on Pattern Recognition (ICPR)*, pages 93–97, 1994. 16

[16] Yuval Eldar, Michael Lindenbaum, Moshe Porat, and Yehoshua Y. Zeevi. The farthest point strategy for progressive image sampling. *Transactions on Image Processing (TIP)*, 6(9):1305–1315, 1997. 16

[17] Zicong Fan, Adrian Spurr, Muhammed Kocabas, Siyu Tang, Michael J. Black, and Otmar Hilliges. Learning to disambiguate strongly interacting hands via probabilistic per-pixel part segmentation. In *International Conference on 3D Vision (3DV)*, pages 1–10, 2021. 2, 6, 18

[18] Thomas Feix, Javier Romero, Heinz-Bodo Schmiedmayer, Aaron M. Dollar, and Danica Kragic. The grasp taxonomy of human grasp types. *Transactions on Human-Machine Systems (THMS)*, 46(1):66–77, 2016. 3

[19] Guillermo Garcia-Hernando, Shanxin Yuan, Seungryul Baek, and Tae-Kyun Kim. First-person hand action benchmark with RGB-D videos and 3D hand pose annotations. In *Computer Vision and Pattern Recognition (CVPR)*, pages 409–419, 2018. 2, 3

[20] Patrick Grady, Chengcheng Tang, Samarth Brahmbhatt, Christopher D. Twigg, Chengde Wan, James Hays, and Charles C. Kemp. PressureVision: Estimating hand pressure from a single RGB image. *European Conference on Computer Vision (ECCV)*, 13666:328–345, 2022. 3

[21] Patrick Grady, Chengcheng Tang, Christopher D. Twigg, Minh Vo, Samarth Brahmbhatt, and Charles C. Kemp. ContactOpt: Optimizing contact to improve grasps. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1471–1481, 2021. 2, 3, 7, 19

[22] Shreyas Hampali, Mahdi Rad, Markus Oberweger, and Vincent Lepetit. HOnnotate: A method for 3D annotation of hand and object poses. In *Computer Vision and Pattern Recognition (CVPR)*, pages 3193–3203, 2020. 1, 2, 3, 4

[23] Shreyas Hampali, Sayan Deb Sarkar, Mahdi Rad, and Vincent Lepetit. Keypoint transformer: Solving joint identification in challenging hands and object interactions for accurate 3D pose estimation. In *Computer Vision and Pattern Recognition (CVPR)*, pages 11090–11100, 2022. 1, 2, 3, 4

[24] Yana Hasson, Bugra Tekin, Federica Bogo, Ivan Laptev, Marc Pollefeys, and Cordelia Schmid. Leveraging photometric consistency over time for sparsely supervised hand-object reconstruction. In *Computer Vision and Pattern Recognition (CVPR)*, pages 568–577, 2020. 2, 7, 19

[25] Yana Hasson, Gül Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael J. Black, Ivan Laptev, and Cordelia Schmid.

Learning joint reconstruction of hands and manipulated objects. In *Computer Vision and Pattern Recognition (CVPR)*, pages 11807–11816, 2019. 1, 2, 3, 7, 19

[26] Chun-Hao P. Huang, Hongwei Yi, Markus Höschle, Matvey Safroshkin, Tsvetelina Alexiadis, Senya Polikovsky, Daniel Scharstein, and Michael J. Black. Capturing and inferring dense full-body human-scene contact. In *Computer Vision and Pattern Recognition (CVPR)*, pages 13274–13285, June 2022. 2, 3

[27] Yinghao Huang, Omid Taheri, Michael J. Black, and Dimitrios Tzionas. InterCap: Joint markerless 3D tracking of humans and objects in interaction. In *German Conference on Pattern Recognition (GCPR)*, volume 13485, pages 281–299, 2022. 2

[28] Umar Iqbal, Pavlo Molchanov, Thomas Breuel Juergen Gall, and Jan Kautz. Hand pose estimation via latent 2.5D heatmap regression. In *European Conference on Computer Vision (ECCV)*, pages 118–134, 2018. 2

[29] Noriko Kamakura, Michiko Matsuo, Harumi Ishii, Fumiko Mitsuboshi, and Yoriko Miura. Patterns of static prehension in normal hands. *American Journal of Occupational Therapy*, 34(7):437–445, 1980. 3

[30] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *Computer Vision and Pattern Recognition (CVPR)*, pages 7122–7131, 2018. 7, 15, 16

[31] 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)*, pages 333–344, 2020. 3, 8

[32] Manuel Kaufmann, Velko Vechev, and Dario Mylonopoulos. aitviewer, 7 2022. 7

[33] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 15

[34] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. VIBE: Video inference for human body pose and shape estimation. In *Computer Vision and Pattern Recognition (CVPR)*, pages 5253–5263, 2020. 2, 5, 16, 17, 18

[35] Muhammed Kocabas, Chun-Hao P. Huang, Otmar Hilliges, and Michael J. Black. PARE: Part attention regressor for 3D human body estimation. In *International Conference on Computer Vision (ICCV)*, pages 11127–11137, 2021. 7, 15, 16

[36] Muhammed Kocabas, Chun-Hao P. Huang, Joachim Tesch, Lea Müller, Otmar Hilliges, and Michael J. Black. SPEC: Seeing people in the wild with an estimated camera. In *International Conference on Computer Vision (ICCV)*, pages 11035–11045, 2021. 15, 16

[37] Taein Kwon, Bugra Tekin, Jan Stühmer, Federica Bogo, and Marc Pollefeys. H2O: Two hands manipulating objects for first person interaction recognition. In *International Conference on Computer Vision (ICCV)*, pages 10138–10148, 2021. 1, 2, 3, 4

[38] Mengcheng Li, Liang An, Hongwen Zhang, Lianpeng Wu, Feng Chen, Tao Yu, and Yebin Liu. Interacting attention graph for single image two-hand reconstruction. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2761–2770, 2022. 2

[39] Xiaolong Li, He Wang, Li Yi, Leonidas J. Guibas, A. Lynn Abbott, and Shuran Song. Category-level articulated object pose estimation. In *Computer Vision and Pattern Recognition (CVPR)*, pages 3703–3712, 2020. 8

[40] 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)*, pages 14687–14697, 2021. 2

[41] Yunze Liu, Yun Liu, Che Jiang, Kangbo Lyu, Weikang Wan, Hao Shen, Boqiang Liang, Zhoujie Fu, He Wang, and Li Yi. HOI4D: A 4D egocentric dataset for category-level human-object interaction. In *Computer Vision and Pattern Recognition (CVPR)*, pages 21013–21022, 2022. 1, 2, 3, 4, 13

[42] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *International Conference on Computer Vision (ICCV)*, pages 5441–5450, 2019. 2, 4, 5

[43] Frank Michel, Alexander Krull, Eric Brachmann, Michael Ying Yang, Stefan Gumhold, and Carsten Rother. Pose estimation of kinematic chain instances via object coordinate regression. In *British Machine Vision Conference (BMVC)*, pages 181.1–181.11, 2015. 13

[44] Gyeongsik Moon, Shoou-I Yu, He Wen, Takaaki Shiratori, and Kyoung Mu Lee. InterHand2.6M: A dataset and baseline for 3D interacting hand pose estimation from a single RGB image. In *European Conference on Computer Vision (ECCV)*, volume 12365, pages 548–564, 2020. 2, 4, 6, 18

[45] Franziska Mueller, Florian Bernard, Oleksandr Sotnychenko, Dushyant Mehta, Srinath Sridhar, Dan Casas, and Christian Theobalt. GANerated hands for real-time 3D hand tracking from monocular RGB. In *Computer Vision and Pattern Recognition (CVPR)*, pages 49–59, 2018. 2

[46] Franziska Mueller, Dushyant Mehta, Oleksandr Sotnychenko, Srinath Sridhar, Dan Casas, and Christian Theobalt. Real-time hand tracking under occlusion from an egocentric RGB-D sensor. In *International Conference on Computer Vision (ICCV)*, pages 1163–1172, 2017. 2

[47] Supreeth Narasimhaswamy, Trung Nguyen, and Minh Hoai Nguyen. Detecting hands and recognizing physical contact in the wild. In *Conference on Neural Information Processing Systems (NeurIPS)*, volume 33, 2020. 3

[48] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *Computer Vision and Pattern Recognition (CVPR)*, pages 10975–10985, 2019. 3, 4, 5

[49] Tu-Hoa Pham, Nikolaos Kyriazis, Antonis A. Argyros, and Abderrahmane Kheddar. Hand-object contact force estimation from markerless visual tracking. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 40(12):2883–2896, 2018. 3

[50] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *Computer Vision and Pattern Recognition (CVPR)*, pages 77–85, 2017. 8

[51] James M. Rehg and Takeo Kanade. Visual tracking of high DOF articulated structures: An application to human hand tracking. In *European Conference on Computer Vision (ECCV)*, volume 801, pages 35–46, 1994. 2

[52] Grégory Rogez, James Steven Supančič III, and Deva Ramanan. Understanding everyday hands in action from RGB-D images. In *International Conference on Computer Vision (ICCV)*, pages 3889–3897, 2015. 3

[53] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *Transactions on Graphics (TOG)*, 36(6):245:1–245:17, 2017. 2, 6

[54] István Sárándi, Timm Linder, Kai O. Arras, and Bastian Leibe. Metric-scale truncation-robust heatmaps for 3D human pose estimation. In *International Conference on Automatic Face & Gesture Recognition (FG)*, pages 407–414, 2020. 7, 16

[55] Fadime Sener, Dibyadip Chatterjee, Daniel Shelepov, Kun He, Dipika Singhania, Robert Wang, and Angela Yao. Assembly101: A large-scale multi-view video dataset for understanding procedural activities. In *Computer Vision and Pattern Recognition (CVPR)*, pages 21064–21074, 2022. 2

[56] Dandan Shan, Jiaqi Geng, Michelle Shu, and David F. Fouhey. Understanding human hands in contact at internet scale. In *Computer Vision and Pattern Recognition (CVPR)*, pages 9866–9875, 2020. 3

[57] Tomas Simon, Hanbyul Joo, Iain Matthews, and Yaser Sheikh. Hand keypoint detection in single images using multiview bootstrapping. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4645–4653, 2017. 2

[58] Adrian Spurr, Aneesh Dahiya, Xi Wang, Xucong Zhang, and Otmar Hilliges. Self-supervised 3D hand pose estimation from monocular RGB via contrastive learning. In *International Conference on Computer Vision (ICCV)*, pages 11210–11219, 2021. 2

[59] Adrian Spurr, Umar Iqbal, Pavlo Molchanov, Otmar Hilliges, and Jan Kautz. Weakly supervised 3D hand pose estimation via biomechanical constraints. In *European Conference on Computer Vision (ECCV)*, volume 12362, pages 211–228, 2020. 2

[60] Adrian Spurr, Jie Song, Seonwook Park, and Otmar Hilliges. Cross-modal deep variational hand pose estimation. In *Computer Vision and Pattern Recognition (CVPR)*, pages 89–98, 2018. 2

[61] Srinath Sridhar, Franziska Mueller, Michael Zollhoefer, Dan Casas, Antti Oulasvirta, and Christian Theobalt. Real-time joint tracking of a hand manipulating an object from RGB-D input. In *European Conference on Computer Vision (ECCV)*, volume 9906, pages 294–310, 2016. 2, 3

[62] Stefan Stevšič and Otmar Hilliges. Spatial attention improves iterative 6D object pose estimation. In *International Conference on 3D Vision (3DV)*, pages 1070–1078, 2020. 6

[63] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Computer Vision and Pattern Recognition (CVPR)*, 2019. 15

[64] Omid Taheri, Vassileios Choutas, Michael J. Black, and Dimitrios Tzionas. GOAL: Generating 4D whole-body motion for hand-object grasping. In *Computer Vision and Pattern Recognition (CVPR)*, pages 13253–13263, 2022. 8, 19

[65] 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)*, volume 12349, pages 581–600, 2020. 2, 3, 4, 5, 14, 15, 19

[66] 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)*, pages 4511–4520, 2019. 2

[67] 3dMDhand system series. https://3dmd.com/products/. 4, 13

[68] Shashank Tripathi, Lea Müller, Chun-Hao P. Huang, Taheri Omid, Michael J. Black, and Dimitrios Tzionas. 3D human pose estimation via intuitive physics. In *Computer Vision and Pattern Recognition (CVPR)*, June 2023. 3

[69] Aggeliki Tsoli and Antonis A. Argyros. Joint 3D tracking of a deformable object in interaction with a hand. In *European Conference on Computer Vision (ECCV)*, volume 11218, pages 504–520, 2018. 3

[70] Dimitrios Tzionas, Luca Ballan, Abhilash Srikantha, Pablo Aponte, Marc Pollefeys, and Juergen Gall. Capturing hands in action using discriminative salient points and physics simulation. *International Journal of Computer Vision (IJCV)*, 118(2):172–193, 2016. 2

[71] Dimitrios Tzionas and Juergen Gall. A comparison of directional distances for hand pose estimation. In *German Conference on Pattern Recognition (GCPR)*, volume 8142, pages 131–141, 2013. 2

[72] Dimitrios Tzionas and Juergen Gall. 3D object reconstruction from hand-object interactions. In *International Conference on Computer Vision (ICCV)*, pages 729–737, 2015. 3

[73] Dimitrios Tzionas and Juergen Gall. Reconstructing articulated rigged models from RGB-D videos. In *European Conference on Computer Vision Workshops (ECCVw)*, volume 9915, pages 620–633, 2016. 2, 3

[74] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research (JMLR)*, 9(86):2579–2605, 2008. 4

[75] Vicon Vantage: Cutting edge, flagship camera with intelligent feedback and resolution. https://www.vicon.com/hardware/cameras/vantage. 2, 4

[76] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, and Hao Su. SAPIEN: A simulated part-based interactive environment. In *Computer Vision and Pattern Recognition (CVPR)*, pages 11094–11104, 2020. 13

[77] Lixin Yang, Xinyu Zhan, Kailin Li, Wenqiang Xu, Jiefeng Li, and Cewu Lu. CPF: Learning a contact potential field to model the hand-object interaction. In *International Conference on Computer Vision (ICCV)*, pages 11097–11106, 2021. 2, 3, 7, 19

[78] Yifei Yin, Chen Guo, Manuel Kaufmann, Juan Zarate, Jie Song, and Otmar Hilliges. Hi4D: 4D instance segmentation of close human interaction. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3

[79] Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. DPOD: 6D pose object detector and refiner. In *International Conference on Computer Vision (ICCV)*, pages 1941–1950, 2019. 6

[80] Baowen Zhang, Yangang Wang, Xiaoming Deng, Yinda Zhang, Ping Tan, Cuixia Ma, and Hongan Wang. Interacting two-hand 3D pose and shape reconstruction from single color image. In *International Conference on Computer Vision (ICCV)*, pages 11354–11363, 2021. 2

[81] He Zhang, Yuting Ye, Takaaki Shiratori, and Taku Komura. ManipNet: Neural manipulation synthesis with a hand-object spatial representation. *Transactions on Graphics (TOG)*, 40(4):1–14, 2021. 8

[82] Hao Zhang, Yuxiao Zhou, Yifei Tian, Jun-Hai Yong, and Feng Xu. Single depth view based real-time reconstruction of hand-object interactions. *Transactions on Graphics (TOG)*, 40(3):29:1–29:12, 2021. 3

[83] Xiong Zhang, Qiang Li, Hong Mo, Wenbo Zhang, and Wen Zheng. End-to-end hand mesh recovery from a monocular RGB image. In *International Conference on Computer Vision (ICCV)*, pages 2354–2364, 2019. 2, 7, 16

[84] Keyang Zhou, Bharat Lal Bhatnagar, Jan Eric Lenssen, and Gerard Pons-Moll. TOCH: Spatio-temporal object-to-hand correspondence for motion refinement. In *European Conference on Computer Vision (ECCV)*, volume 13663, pages 1–19, 2022. 3, 19

[85] 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)*, pages 5745–5753, 2019. 16

[86] Yuxiao Zhou, Marc Habermann, Weipeng Xu, Ikhsanul Habibie, Christian Theobalt, and Feng Xu. Monocular real-time hand shape and motion capture using multi-modal data. In *Computer Vision and Pattern Recognition (CVPR)*, pages 5345–5354, 2020. 2

[87] Andrea Ziani, Zicong Fan, Muhammed Kocabas, Sammy Christen, and Otmar Hilliges. TempCLR: Reconstructing hands via time-coherent contrastive learning. In *International Conference on 3D Vision (3DV)*, pages 627–636, 2022. 2, 18

[88] Christian Zimmermann and Thomas Brox. Learning to estimate 3D hand pose from single RGB images. In *International Conference on Computer Vision (ICCV)*, pages 4913–4921, 2017. 2

[89] Christian Zimmermann, Duygu Ceylan, Jimei Yang, Bryan Russell, Max Argus, and Thomas Brox. FreiHAND: A dataset for markerless capture of hand pose and shape from single RGB images. In *International Conference on Computer Vision (ICCV)*, pages 813–822, 2019. 2, 3

Appendix - Supplementary Material

1. Dataset Details

Objects in ARCTIC: Figure 1 shows all 11 articulated objects in our dataset. Objects in ARCTIC consists of two rigid parts that rotate about an axis. Each dash line in the figure shows the articulation axis.

Marker sets: Figure 2 shows marker sets for an object, the full human body, two hands, and the egocentric camera along with the marker size. Markers in this visualization are shown to scale. The marker locations for all objects and all subjects can be found in the data release.

Dataset statistics: Table 1 shows the number of images and the number of sequences per subject for our dataset. The average sequence length is 698 frames (view-agnostic), corresponding to 23.3 seconds. In total, we have 2.1M images and there are more than 200k images for most subjects. Table 2 shows the number of images per object. All objects have more than 170k images. To encourage different modes of interaction, we capture different intents for each object: "use" and "grasp". Although both are for dexterous manipulation, in the "use" sequences, the subjects are allowed to articulate the object but not in the "grasp". Since we focus on studying articulation, we capture more "use" sequences.

Protocol splits: Table 3 shows the number of images and subjects in the allocentric and the egocentric settings. Both settings use the same subject split – 8 subjects for training, 1 for validation and 1 for testing. The allocentric setting uses images from the 8 allocentric static views for training, validation, and testing. The egocentric setting, in the training split, we allow models to use images from all 9 views for additional supervision; During inference, however, models are evaluated with only egocentric images.

Depth images: Since we perform full-body capture, we can render depth images with full-body interaction. Since most existing articulated object datasets contain neither two-hands nor human bodies [41, 43, 76], and having a human in the scene is a realistic setting, we believe that ARCTIC brings additional challenges of heavy occlusion and dynamic manipulation to the depth community. Figure 3 shows examples of the depth images and the corresponding RGB images. The depth can be rendered with any synthetic sensor noise model (e.g., Kinect, right).

Subjects	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	Total
# Images	209k	224k	220k	212k	227k	228k	191k	280k	208k	135k	2.1M
# Seqs	34	37	38	31	34	36	29	42	37	21	339

Table 1. **Number of images and sequences for each subject**. The average sequence length is 698 frames (view-agnostic), corresponding to 23.3 seconds.

2. Data Capture Details

2.1. MoCap System with 54 Cameras

Figure 4 illustrates our MoCap system. When capturing quality hand-object interaction data, the key is to eliminate occlusion in hand self-occlusion, hand-hand occlusion, and hand-object occlusion settings. To minimize these sources of occlusion, we use 54 high-resolution MoCap cameras during our capture. The camera positions and orientations are shown in Fig. 4a in a side view and a bird-eye view. We also show the markers tracked by our system in Fig. 4b. The system tracks markers on the egocentric camera, the human subject (full body with hands), the articulated object (notebook in this case), and props such as the table.

2.2. Creating Personalized Template

To create a personalized full body-and-hand template mesh for each subject, we obtain 3D scans of the subjects in varying poses using a 3dMD scanner [67]. We then register the SMPL-X model to the scans to obtain aligned meshes. The registered SMPL-X meshes are unposed to a canonical T-Pose. We perform a SMPL-X model-based fitting to the unposed meshes using vertex-to-vertex distances with the SMPL-X vertex correspondences. Fitting with multiple T-Posed meshes allows us to filter out potential noise, and to capture the occluded regions of the body and hands, resulting in a reliable personalized SMPL-X template for each subject.

2.3. Estimating Rotation Axis and Articulated Pose

To solve for the rotation axis of each articulated object, we attach markers on each rigid part of each object and capture a calibration MoCap sequence by articulating its two parts. Since the two parts rotate about an axis, the trajectory of each marker follows a circle on a 2D plane. We solve for the center of each circle using least-squares, and fit a 3D line through the centers to obtain an initial rotation axis estimate. We then refine the rotation axis estimate by minimizing a cost function.

Formally, let $X_t^i \in \mathbb{R}^3$ denote the 3D position of a marker i at time step t placed on the "top" part of the object (*e.g.*, the lid of the ketchup bottle); $Y_t^j \in \mathbb{R}^3$ denotes a marker j on the "bottom" part of the object (*e.g.*, the main



Figure 1. **ARCTIC objects**. Each line shows the articulation axis.

Objects	Notebook	Box	Espresso machine	Waffle iron	Laptop	Phone	Capsule machine	Mixer	Ketchup bottle	Scissors	Microwave	Total
Use	163k	152k	141k	166k	171k	159k	159k	156k	138k	128k	144k	1.7M
Grasp	27k	36k	46k	41k	40k	49k	37k	47k	51k	44k	43k	0.4M
Total	190k	187k	187k	207k	211k	208k	196k	202k	189k	172k	186k	2.1M

Table 2. **Number of images for each object in ARCTIC**. In ARCTIC, we focus on studying hand interaction with object articulation. Therefore, we capture more "use" sequences in which subjects can articulate the object. To encourage different modes of interaction, we also ask subjects to "grasp" the objects without articulating the objects. Since we focus on object articulation, we capture more data for "use" than in "grasp".

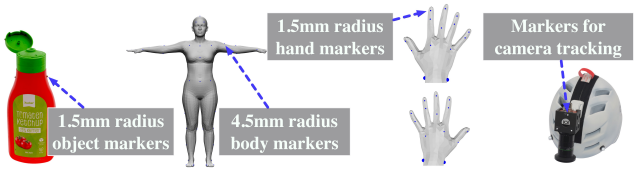


Figure 2. **Markers for motion capture (MoCap)**. We put 1.5mm radius markers on objects, hands and the egocentric camera. For the body, we use 4.5mm radius markers. The markers are shown here to scale. Best viewed in color and zoomed-in.

Splits	# Train Images	# Val Images	# Test Images
allo	1.5M	202k	195k
ego	1.7M	25k	24k

Table 3. **Number of images for each protocol**.

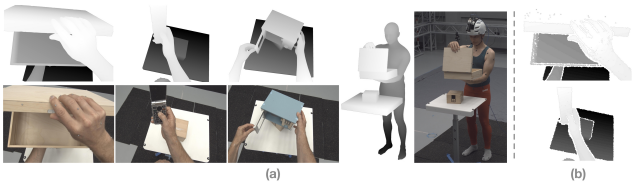


Figure 3. **Rendered depth images of human-object interaction.** (a) Rendered depth images of our 3D data, and the corresponding RGB images, and (b) depth images with synthetic kinect noise. Best viewed zoomed-in. For better depth-map visualization we do not render the floor here.

body of the ketchup bottle) at time t. We pick a frame t_0 corresponding to a pre-defined rest pose for each object (for example, a frame in which the lid of the ketchup bottle is closed). We then transform the MoCap sequence $\{X_t^i\}_{i,t}$ and $\{Y_t^j\}_{j,t}$ into a canonical sequence $\{\bar{X}_t^i\}_{i,t}$ and $\{\bar{Y}_t^j\}_{j,t}$ via a rigid transformation (R, T) such that $\bar{Y}_t^j = Y_{t_0}^j$ for all t and j. After the canonicalization, the markers on the bottom part of the object are stationary across the entire sequence, and the top markers rotate around an axis. Further, the trajectory of each marker is a circle on a 2D plane. We fit a circle to the trajectory $\{\bar{X}_t^i\}_{t=1,\cdots,N}$ of an arbitrary marker i using least-squares, and convert the 2D circle center to the 3D space (x_i, y_i, z_i). We then fit a 3D line (v, v_0) to the center of each circle $\{(x_i, y_i, z_i)\}_i$ in 3D, where $v \in \mathbb{R}^3$ is a unit directional vector for the line and $v_0 \in \mathbb{R}^3$ is an arbitrary anchor point that the line crosses. Since fitting a 3D line to 3D centers can be imprecise, we refine the rotation axis further by minimizing the following cost function

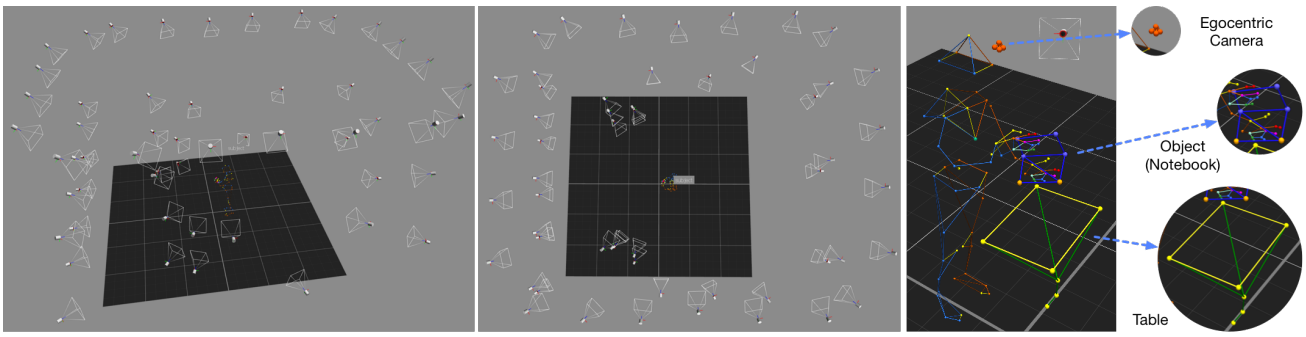
$$(v^*, v_0^*, \omega^*_{t=1,\cdots,N}) = \underset{v,v_0,\omega_{t=1,\cdots,N}}{\arg\min} \sum_i ||\bar{X}_t^i - f(\bar{X}_{t_0}^i | v, v_0, \omega_t)||_2^2 \tag{8}$$

where N is the number of frames in a sequence; $\omega^*_{t=1,\cdots,N}$ are the articulation angles in radians for all the frames relative to the rest pose. The quantities v^*, v_0^* are the refined rotation axis for the object. The function $f(\bar{X}_{t_0}^i | v, v_0, \omega_t)$ rotates $\{\bar{X}_t^i\}$ about the estimated axis (v^*, v_0^*) by ω_t, the amount of articulation. This ensures that the estimated rotation axis is consistent with the marker trajectory in the MoCap data.

To define the articulated object pose, we need to estimate the 1D articulation angle, and its 6D rigid pose. To compute the former, after the rotation axis (v^*, v_0^*) is estimated, for an actual MoCap sequence, the articulation angles are obtained by performing a 2D projection of the 3D marker positions along the rotation axis. Since the 2D projection lies on a circle, the articulation angle can be estimated arithmetically. The articulation angle is measured relative to each object's rest pose defined in t_0 during the rotation axis estimation step. We take the median of the articulation angles estimated from all markers at a time step as our ground-truth articulation angle. Finally, to define the articulated object pose, we also need the 6D object pose for its orientation and translation. To solve for the 6D pose (R_t, T_t) for a frame t, we compute the rigid transformation from $\bar{Y}_{t_0}$ to Y_t. In other words, we compute the 6D pose using the base marker according to its correspondence from the canonical space to the MoCap space.

2.4. Computing Hand-Object Binary Contact

We consider the two hands and the two parts of each articulated object as four watertight meshes for computing ground-truth binary contact labels. Given one mesh from the hands and one mesh from the object parts, we follow GRAB [65] to compute vertex-level contact. The main idea in GRAB is to label vertices on a mesh as in contact with another mesh based on two cases: "contact under-shooting" and "contact over-shooting". When vertices on a mesh are



(a) MoCap system with 54 cameras (left: side view; right: bird-eye view) (b) Observed markers during data capture

Figure 4. **Our capture system with 54 high-resolution MoCap cameras**. (a) MoCap system in a side view and a bird-eye view, illustrating the 54 MoCap cameras used to eliminate occlusion during the capture. (b) Observed markers for a captured frame, showing markers tracking the full human body with hands, the object (notebook in this case), the egocentric camera, and props such as the table. Best viewed in color and zoomed in.

Hand Branch		
Nr.	**Module**	**Details**
1	pool	AvgPool2d(output_size=1)
2	cam_init	Linear(in_dim=2048, out_dim=512, bias=True)
3	cam_init	ReLU()
4	cam_init	Linear(in_dim=512, out_dim=512, bias=True)
5	cam_init	ReLU()
6	cam_init	Linear(in_dim=512, out_dim=3, bias=True)
7	refine.fwd	Concat(["feat", "hand_pose", "cam", "shape"])
8	refine.fwd	Linear(in_dim=2157, out_dim=1024, bias=True)
9	refine.fwd	ReLU()
10	refine.fwd	Dropout(p=0.5)
11	refine.fwd	Linear(in_dim=1024, out_dim=1024, bias=True)
12	refine.fwd	ReLU()
13	refine.fwd	Dropout(p=0.5)
14	refine.decode.pose_6d	Linear(in_dim=1024, out_dim=96, bias=True)
15	refine.decode.shape	Linear(in_dim=1024, out_dim=10, bias=True)
16	refine.decode.cam	Linear(in_dim=1024, out_dim=3, bias=True)
Object Branch		
1	pool	AvgPool2d(output_size=1)
2	cam_init	Linear(in_dim=2048, out_dim=512, bias=True)
3	cam_init	ReLU()
4	cam_init	Linear(in_dim=512, out_dim=512, bias=True)
5	cam_init	ReLU()
6	cam_init	Linear(in_dim=512, out_dim=3, bias=True)
7	refine.fwd	Concat(["feat", "rot", "cam", "arti"])
8	refine.fwd	Linear(in_dim=2055, out_dim=1024, bias=True)
9	refine.fwd	ReLU()
10	refine.fwd	Dropout(p=0.5)
11	refine.fwd	Linear(in_dim=1024, out_dim=1024, bias=True)
12	refine.fwd	ReLU()
13	refine.fwd	Dropout(p=0.5)
14	refine.decode.rot	Linear(in_dim=1024, out_dim=3, bias=True)
15	refine.decode.cam	Linear(in_dim=1024, out_dim=3, bias=True)
16	refine.decode.arti	Linear(in_dim=1024, out_dim=1, bias=True)

Table 4. **Details of the decoder in ArcticNet-SF.**

not inside another mesh, it is considered "under-shooting", and geometric proximity is used to label contact. When there is interpenetration between two meshes, for example, the thumb goes through a thin structure, the vertices of the thumb that "over-shoot" the thin structure are labeled as in contact as well as the vertices that are inside the structure. For more details, we refer readers to [65].

3. Model Details and Results

General implementation details: For all experiments, we use a ResNet-50 [63] backbone pre-trained on ImageNet [13]. The models are trained with the Adam optimizer [33] using a learning rate of $1e^{-5}$. For visibility, we crop each image around a square region centered around the object and resize the image to 224×224. Data augmentation is applied to the input image: rotation ($\pm 30°$), scaling ($\pm 25\%$), and color jittering ($\pm 40\%$).

3.1. ArcticNet

ArcticNet-SF Architecture: We show the details of ArcticNet-SF in Table 4. ArcticNet-SF uses an encoder-decoder architecture. Given an input image, we use average pooling to obtain a single image feature vector with dimension 2048 from the backbone. The image feature vector is used to predict the initial camera parameters (in "cam_init"). Following [30], we use an iterative refinement scheme to predict parameters of the hands and the objects. First, we initialize all parameters to zero except for the camera parameters, which we initialize with the prediction in "cam_init". For the hand branch, we concatenate the image feature vector, the initial hand poses, the hand shape vector into a single vector, and the predicted camera parameters for refinement (Line 7-16). Within the refinement step, we first predict a latent vector using an MLP (Line 7-13). The latent vector is then being decoded via different heads to the residuals for MANO joint angles, shape and camera parameters (Line 14-16). The decoded residuals are added to the current estimates of the parameters respectively and will be used as inputs for the next refinement step (Line 7-16 again). We have two iterations for the refinement. The object branch has a similar refinement scheme, but instead it predicts the object rotation, camera parameters, and object articulation. Following [35, 36], we

use the 6D rotation representation in [85] for MANO joint angles. Following [5, 30, 35, 54, 83], for the predicted weak perspective camera parameters, we use a fixed focal length of 1000.0 and convert them to translations for each entity in the scene.

ArcticNet-LSTM Architecture: The LSTM model takes in a moving window of images and estimates 3D meshes for each frame. We use the same structure as ArcticNet-SF except that the image features within each window are passed through an LSTM network before being decoded. We use a bidirectional LSTM with two hidden layers, hidden dimension of 1024, and a window size of 11 based on validation.

Training losses: For each frame, our loss $\mathcal{L}$ is defined as the sum of the left hand, right hand, object, and interaction losses: $\mathcal{L} = \mathcal{L}_l + \mathcal{L}_r + \mathcal{L}_o + \mathcal{L}_{int}$. In particular, the hand losses are defined as

$$\mathcal{L}_h = \lambda^h_{3D}\mathcal{L}^h_{3D} + \lambda^h_{2D}\mathcal{L}^h_{2D} + \lambda^h_{\Theta}\mathcal{L}^h_{\Theta} + \lambda^h_T\mathcal{L}^h_T, \quad (9)$$

where $h = \{l, r\}$ denotes the handedness. We fully supervise the 3D joints (after subtracting the roots), the 2D reprojection of the predicted 3D joints, the MANO pose and shape parameters and the weak-perspective camera parameters. Similarly, we pre-define 3D landmarks for objects using farthest point sampling [15, 16] on the object mesh. Using these landmarks, we formulate the object losses as

$$\mathcal{L}_o = \lambda^o_{3D}\mathcal{L}^o_{3D} + \lambda^o_{2D}\mathcal{L}^o_{2D} + \lambda_\omega\mathcal{L}_\omega + \lambda_R\mathcal{L}_R + \lambda^o_T\mathcal{L}^o_T, \quad (10)$$

where $\mathcal{L}_\omega$, $\mathcal{L}_R$ and $\mathcal{L}^o_T$ supervise the articulation angle in radians, the global orientation and the weak perspective camera parameters. For the interaction loss $\mathcal{L}_{int}$, we use the contact deviation (CDev) metric (see main paper) as a loss term to improve hand-object contact. We apply this loss between the left-hand/object, and right-hand/object. The loss $\mathcal{L}_{int}$ is a sum of the two. All losses above use the MSE criterion. All λ variables are hyper-parameters and are set empirically based on validation performance. In particular, we set all λs to 1.0 except $\lambda^*_{3D} = 5.0$, $\lambda^h_{2D} = 5.0$, $\lambda^h_{\Theta} = 10.0$, $\lambda^h_{\beta} = 0.001$ where $*$ denote a hand or an object.

Training details: We train with a batch size of 64. For the allocentric setting, we train single-frame models for 20 epochs. Since training temporal model is computation intensive, following VIBE [34], we dump image features of pre-trained single-frame models to disk then train the LSTM models directly on the image features for 10 epochs. For the egocentric setting, since a model has access to both allocentric and egocentric images during training, to speed up training, we finetune pre-trained allocentric models on egocentric training images (1 camera) for 50 epochs.

Camera model: Following previous work on body and hand surface reconstruction [5, 30, 35, 36, 54, 83], to estimate the translation of hands and objects (T_l, T_r, T_o), we predict weak-perspective camera parameters (s, t_x, t_y) for each entity in the scene. The camera parameters consist of the scale $s \in \mathbb{R}$ and translation $(t_x, t_y) \in \mathbb{R}^2$ in pixel space and the translation can be recovered from (s, t_x, t_y) [35, 36] via:

$$T = (t_x, t_y, \frac{2f}{ws}) \in \mathbb{R}^3. \quad (11)$$

The terms w and f are the patch size and the focal length. We do this for each (T_l, T_r, T_o).

Qualitative Results: Figure 5 shows the predictions of ArcticNet-SF and ArcticNet-LSTM on the test set. As shown in the quantitative results in the main manuscript, the ArcticNet-LSTM model has lower errors overall for its prediction and it has better contact. This is consistent with the observations in the qualitative examples here. We hypothesize that this is because the LSTM allows the network to jointly reason between the motions of hands and objects.

3.2. InterField

InterField-SF Architecture: Table 5 details our InterField model. As an example, we illustrate how the right hand interaction field is predicted. The left hand, and the object are predicted in a similar way. In particular, from an input image, we obtain a 2048-dimensional image feature vector from the image backbone. The vector is passed through an MLP and is projected to lower dimension for computational efficiency (Line 1-4). We use subsampled vertices of the hand, and concatenate the 3D location of each vertex of the subsampled hand in the canonical pose with the 512-dimensional image feature vector, resulting a point cloud with 515 dimensions. The point cloud is passed through a PointNet backbone to obtain a latent point cloud with 512 dimensions (Line 5-11). Within the PointNet backbone, the 515-dimensional input point cloud is passed through a sequence of layers to produce lower level point features (Line 5-6). The point features are further processed through Line 7-11. We then concatenate the point cloud from the shallow layers (output of Line 6) and the deeper layers (output of Line 11) along the feature dimension, resulting in a point cloud whose individual points are in 1024-dimensional. A regressor maps each point (1024-dimensional) to a single scalar for distance prediction (Line 12-18). Finally, we upsample the subsampled distances to the full hand mesh (Line 19). We predict the interaction field of the left hand and the object in the same way. All entities shared the same image and PointNet backbones.

InterField-LSTM Architecture: The LSTM model takes in images from a window and estimates the interaction field for each frame. In particular, we use the same architecture as in ArcticNet-SF except that we pass the image features in a window to an LSTM network before regressing the distances. We use a bidirectional LSTM with two hidden layers, hidden dimension of 1024 and a window size of 11 based on validation performance.

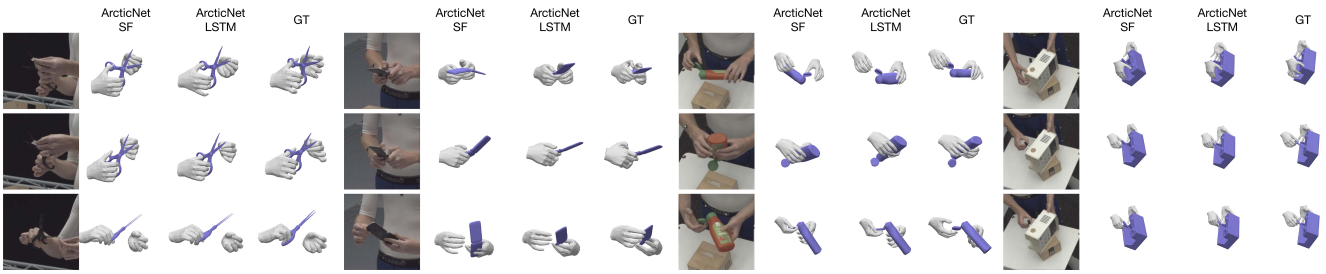
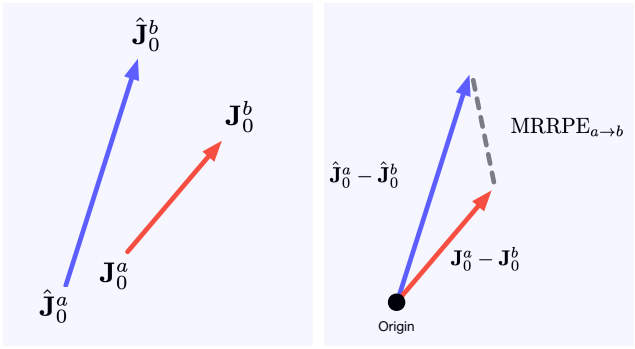


Figure 5. **Qualitative results of ArcticNet-SF and ArcticNet-LSTM**. Best viewed in color and zoomed in.



(a) Predicted and ground-truth roots of entities *a* and *b*

(b) MRRPE for entities *a* and *b*

Figure 6. **An illustration of the MRRPE$_{a\to b}$ metric**. (a) The predicted roots of entities a and b are denoted by $\hat{\mathbf{J}}_0^a$ and $\hat{\mathbf{J}}_0^b$, and $\mathbf{J}_0^a$ and $\mathbf{J}_0^b$ are the corresponding ground-truth. (b) Subtract a by b; MRRPE$_{a\to b} \in \mathbb{R}$ is indicated by the dash line.

Nr.	Module	Details
1	img_feat.down	Linear(in_dim=2048, out_dim=512, bias=True)
2	img_feat.down	ReLU()
3	img_feat.down	Linear(in_dim=512, out_dim=512, bias=True)
4	img_feat.down	ReLU()
5	pointnet.shadow	Linear(in_dim=515, out_dim=512, bias=True)
6	pointnet.shadow	BatchNorm1d(512, affine=True)
7	pointnet.deep	Linear(in_dim=515, out_dim=512, bias=True)
8	pointnet.deep	BatchNorm1d(512, affine=True)
9	pointnet.deep	ReLU()
10	pointnet.deep	Linear(in_dim=515, out_dim=512, bias=True)
11	pointnet.deep	BatchNorm1d(512, affine=True)
12	regressor	Linear(in_dim=1024, out_dim=512, bias=True)
13	regressor	BatchNorm1d(512, affine=True)
14	regressor	ReLU()
15	regressor	Linear(in_dim=512, out_dim=128, bias=True)
16	regressor	BatchNorm1d(128, affine=True)
17	regressor	ReLU()
18	regressor	Linear(in_dim=128, out_dim=1, bias=True)
19	upsample	Linear(in_dim=195, out_dim=778, bias=True)

Table 5. **Details of InterField-SF architecture.**

Training details: For each frame, the network outputs are $\hat{\boldsymbol{F}}^{l\to o}$, $\hat{\boldsymbol{F}}^{r\to o}$, $\hat{\boldsymbol{F}}^{o\to l}$, and $\hat{\boldsymbol{F}}^{o\to r}$. To supervise training, we extract the ground-truth interaction fields for each frame from ARCTIC and formulate an L1 loss $\mathcal{L} = \mathcal{L}_F(l,o) + \mathcal{L}_F(r,o) + \mathcal{L}_F(o,l) + \mathcal{L}_F(o,r)$ where $\mathcal{L}_F(a,b) = ||\boldsymbol{F}^{a\to b} - \hat{\boldsymbol{F}}^{a\to b}||_1$ for entities a and b. For tractability and focus on close interaction, we threshold the interaction field distances at 10cm for training and evaluation.

We train with a batch size of 64 for single-frame models and 32 for LSTM models. For the allocentric setting, we train single-frame models for 20 epochs. Since training temporal model has high computational requirements, following [34], we dump image features of pre-trained single-frame networks to disk and train the LSTM models on the image features for 6 epochs. For the egocentric setting, a model has access to both allocentric and egocentric images in the training set. To speed up training, we finetune pre-trained allocentric models on egocentric images (1 camera) for 100 epochs and 50 epochs for single-frame and LSTM models respectively.

Qualitative Results: Figure 9 shows the predictions of the single-frame model, and the corresponding ground-truth. We use ground-truth hand and object poses for visualization purposes. They are not the inputs of our network. Here we focus on the colors on the meshes; Brighter colors represent closer distances in the interaction fields. The figure shows the feasibility of the task because the predictions correlate well with the ground-truth.

4. Metrics and Experiments

4.1. Metric Details

Acceleration Error (ACC): Following [34], we report acceleration error in m/s^2 to measure the smoothness of consistent motion reconstruction, and interaction field estimation. Formally, suppose $\hat{\mathbf{h}}_i^t \in \mathbb{R}^d$ is the predicted vertex (or distance value) i at frame t of a hand; $\mathbf{h}_i^t$ is the corresponding ground-truth. We compute the corresponding acceleration vector $\hat{\mathbf{u}}_i^t \in \mathbb{R}^d$ of $\hat{\mathbf{h}}_i^t$. Similarly, we compute the acceleration vector $\mathbf{u}_i^t$ for the ground-truth. The acceleration error for a hand is computed as:

$$\frac{1}{TV_h} \sum_{t=1}^{T} \sum_{i=1}^{V_h} \left\| \hat{\mathbf{u}}_i^t - \mathbf{u}_i^t \right\| \tag{12}$$

where V_h, T, d are the number of hand vertices, the sequence length, and the number of dimension for the predic-

	Contact and Relative Position		Motion		Hand	Object	
Object	CDev$_{ho}$ $[mm]\downarrow$	MRRPE$_{rl/ro}$ $[mm]\downarrow$	MDev$_{ho}$ $[mm]\downarrow$	ACC$_{h/o}$ $[m/s^2]\downarrow$	MPJPE$_h$ $[mm]\downarrow$	AAE $[°]\downarrow$	Success Rate $[\%]\uparrow$
Notebook	37.4	47.7/39.8	9.9	5.0/6.5	20.8	3.3	80.4
Box	47.5	66.3/49.2	10.6	5.5/6.7	24.5	1.3	88.2
Espresso machine	48.9	52.5/46.2	9.5	4.8/5.0	24.5	11.0	81.0
Waffle iron	41.8	43.3/39.0	14.6	5.6/7.9	21.3	3.1	74.0
Laptop	42.6	54.7/40.5	12.8	5.2/7.2	21.7	1.7	84.4
Phone	29.5	42.2/31.1	7.5	4.6/7.2	18.8	3.9	62.3
Capsule Machine	30.5	37.6/30.9	7.8	4.7/4.4	19.2	6.9	69.3
Mixer	34.5	41.2/33.9	8.6	4.8/5.3	21.3	2.6	78.3
Ketchup bottle	33.0	45.6/35.0	10.8	5.4/7.4	20.7	7.0	59.2
Scissors	25.6	39.7/22.2	5.8	4.1/5.0	17.7	10.5	50.1
Microwave	60.8	62.6/41.9	9.3	5.2/5.2	26.0	7.3	74.3

Table 6. **Detailed breakdown on test set evaluation per object**. Here we provide the detailed breakdown of the test set evaluation according to each object. For each metric, we use red to denote the object with the highest error; we use blue to denote the lowest error.

	Contact and Relative Position		Motion		Hand	Object	
Size	CDev$_{ho}$	MRRPE$_{rl/ro}$	MDev$_{ho}$	ACC$_{h/o}$	MPJPE$_h$	AAE	Success R.
5	39.9	48.0/37.4	9.3	6.2/8.2	23.3	6.1	72.7
11	**39.0**	**47.0/36.6**	**8.8**	**6.1**/7.7	**22.8**	**5.8**	**74.6**
15	39.7	47.8/36.8	9.0	6.2/**7.6**	22.9	**5.8**	74.4

Table 7. **Effects of window size on ArcticNet-LSTM**. Here we ablate the effect of window size on our model on the validation set.

	Contact and Relative Position		Motion		Hand	Object	
CDev loss	CDev$_{ho}$	MRRPE$_{rl/ro}$	MDev$_{ho}$	ACC$_{h/o}$	MPJPE$_h$	AAE	Success R.
✗	49.0	53.1/45.6	11.9	7.3/10.1	**23.0**	6.1	71.3
✓	**41.9**	**50.1/37.6**	**10.4**	**7.3/9.8**	23.1	**5.9**	**71.8**

Table 8. **Effects of contact deviation (CDev) as a training loss**. Here we ablate the effect of the contact deviation metric as a loss on ArcticNet-SF on the validation set.

tion of a task. To compute the acceleration, we use centered difference:

$$u_i^t = \frac{h_i^{t-1} - 2h_i^t + h_i^{t+1}}{w^2} \tag{13}$$

where $w = 1/30s$ is the stencil width of 30-FPS videos. Note that previous methods [34, 87] computing the acceleration errors did not divide the error by w^2, leading to significantly smaller errors. The prediction dimension (d) for the reconstruction task, and the interaction field task are 3 and 1 respectively. For the former, we use root-relative vertices. We compute the acceleration of the object in the same way.

Average Articulation Error (AAE): Suppose $\omega^t \in \mathbb{R}$ and $\hat{\omega}^t \in \mathbb{R}$ are the predicted and ground-truth object articulation at frame t, the average articulation error is defined as

$$\frac{1}{T}\sum_{t=1}^{T} |\omega^t - \hat{\omega}^t| \tag{14}$$

where T is the number of frames.

Mean Relative-Root Position Error (MRRPE): Following [17, 44], to measure the relative root translation between two entities a and b in the scene (a hand or an object),

$$\text{MRRPE}_{a\to b} = \left\| (\mathbf{J}_0^a - \mathbf{J}_0^b) - \left(\hat{\mathbf{J}}_0^a - \hat{\mathbf{J}}_0^b\right) \right\|_2, \tag{15}$$

where $a \in \{l, r, o\}$ and $b \in \{l, r, o\}$ and l, r, o denote the left hand, right hand, and the object, $\mathbf{J}_0 \in \mathbb{R}^3$ is the ground-truth root joint location and $\hat{\mathbf{J}}_0$ the predicted one. Figure 6 shows a visualization of the metric. Suppose we want to compute MRRPE$_{a\to b}$, and $\hat{\mathbf{J}}_0^a$ and $\hat{\mathbf{J}}_0^b$ denote the predicted roots for entities a and b and the notations without "hat" are the ground-truth. The MRRPE value is the distance indicated in the dash line.

4.2. Ablation and Analysis

Detailed analysis on test set: To see a performance breakdown per object, Table 6 shows the evaluation of the ArcticNet-LSTM model evaluated on the test set of the allocentric split. We use red to denote with the worst value for each metric and blue to denote the best value. We can see that the microwave is the hardest object in terms of hand reconstruction (see MPJPE) and contact (see CDev$_{ho}$). This is because when opening the microwave door, the fingers are often heavily occluded. In contrast, the scissors has the lowest hand reconstruction error because it is a small object. In terms of estimating the articulation angle, the box, however, has the smallest error (see AAE). We hypothesize that the articulation angle is easy to observe because the box is the largest object. In contrast, the espresso machine has a small handle, which can be occluded heavily, so it has the highest AAE error. Finally, the scissors is the hardest object to reconstruct its pose (see success rate) because it is a small object and in some view its dark texture is similar to the background color.

Window size for ArcticNet-LSTM: Table 7 shows the effect of window size on ArcticNet-LSTM in the validation set. With more frames in a window, overall the model has better performance. To balance performance and efficiency, we use a window size of 11 for our model.

Contact deviation (CDev) as training loss: Table 8 shows the effect of the contact deviation metric as a training loss

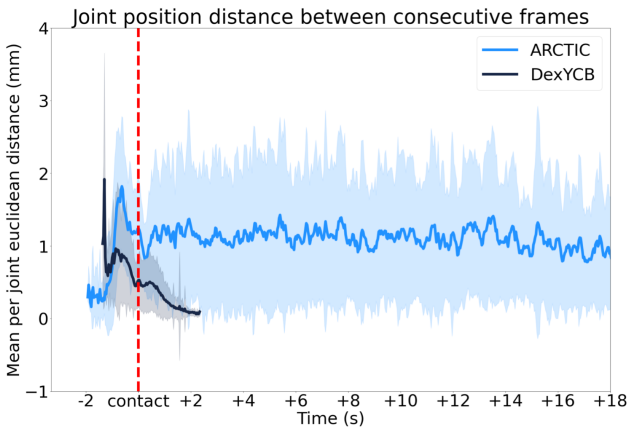


Figure 7. Changes in hand poses after first contact.



Figure 8. **Interaction field estimation on HO3D**

to encourage hand-object contact using the allocentric protocol. The results are evaluated on the validation set. In this ablation, we train the ArcticNet-SF by turning on and off the contact deviation loss. The loss is applied between left-hand/object, and right-hand/object as described in Sec. 3. Results show that the CDev loss improves hand-object contact indicated by the CDev_{ho} metric.

Number of views for ArcticNet-SF: Since ARCTIC has only 1 egocentric view, we ablate the effect of the number of allocentric views. When trained with 2, 4, 6, 8 views, randomly selected, MPJPEs for hands are $49.0mm$, $35.0mm$, $25.8mm$, $23.4mm$; the object success rates are 40.8%, 57.0%, 62.1%, 68.6% on the validation set.

Quantifying dexterous motion: Existing datasets do not show significant changes in hand pose. In particular, poses in ContactPose are fixed relative to the object while DexYCB poses do not change much once contact is established (see Fig. 7). The figure plots the relative change in 3D joints across consecutive frames, without global translation and rotation. The vertical dashed line indicates the first contact. Fig. 2 (main paper) shows that ARCTIC has more diversity in hand poses and contact patterns, resulting from dexterous manipulation, compared to other datasets.

Interaction field estimation on HO3D: Our proposed task can be applied to existing hand datasets with rigid objects. To show this, we trained InterField-SF on HO3D. Fig. 8 shows qualitative results. However, we note that existing hand-object datasets have similar hand poses within each sequence (thus, similar contact) and fewer training images, which are easy to overfit to (see Fig. 8 failure case); ARCTIC is large-scale and it is more challenging due to more dynamic interaction (changing poses and contact) and thus will help in fostering future research.

Improve hand and rigid objects with ARCTIC: We pre-train ArcticNet on ARCTIC, finetune on HO3D (hand + rigid objects). This model is compared to a model trained only on HO3D. Following the HO3D protocol, pre-training on ARCTIC improves MPJPE (scale-translation aligned) errors by 9.2%. For the object, the vertex-to-vertex error (root aligned) improves by 7.1%. This shows that articulated hand-object data benefits hand and rigid object reconstruction.

5. Visualizing ARCTIC

To visualize the our 3D annotation in the dataset, Fig. 10 shows random samples of the 3D meshes of hands and objects overlaid on the images in our ARCTIC dataset. See the video on our project page for our rendered sequences.

6. Discussions and Limitations

We introduce ARCTIC, the first dataset for two hands dexterously manipulating articulated objects and baselines for the task of consistent motion reconstruction, and interaction field estimation. Being a first step, our work is not without limitations.

Known object models in ArcticNet: Similar to existing methods [14, 24, 25, 77], one limitation of our baselines is that they assume known object models. We view articulated 3D shape estimation of unknown objects as an orthogonal problem on which the field is making progress. Now that we have showed the feasibility of inferring hand-object interaction for such objects, future work should bring together our method with 3D articulated object inference. This is challenging and we believe it is critical to make progress on sub-problems, for which ARCTIC can be leveraged.

Toy objects: Some of our objects are toys, which are not to scale and lack some of the visual complexity of real objects. However, the aim of ARCTIC is to study the physical dynamics between hand-object motions.

SMPL-X for capturing contact: Since we use SMPL-X/MANO as our human representation, the human geometry does not capture skin deformation during contact. While a deformable human body/hand model would be ideal for capturing true contact, developing such models is not a goal of ARCTIC. Further, marker and image data ARCTIC can be used to fit a deformable model if developed. Our 3D annotation and contact capture pipeline follows GRAB [65]. In particular, we use MoSh++ to fit SMPL-X to the markers, producing highly accurate fits. We adapt the contact capture pipeline from GRAB. GRAB contact labels are widely used in the community to support projects such as [21, 64, 84].

Markers on hands in our RGB images: We use optical marker-based capture to provide accurate hand and object poses, thus potentially introducing label noise. However, our hand markers are minimally intrusive (1.5mm in radius) and barely visible when images are resized for inputs.

Degree of freedom in our objects: We construct ARCTIC with objects of 1 DoF. This is because many items designed for human interaction often have a single axis of rotation as they are easy to produce and intuitive to manipulate (*e.g.*, doors, refrigerators, ovens, pliers, *etc*.). Thus, objects in ARCTIC are representative of a broad class of objects found in homes and businesses. Importantly, reconstructing interaction with such objects involves occlusion, depth ambiguity, and contact estimation. These issues also apply to objects with more DoF. Further, we capture more diverse hand poses and more challenging hand-object interactions compared to existing hand-object datasets. Future work should expand the number and complexity of objects to further study the problems of depth ambiguity and occlusion.

Approval for human subject data: Subject data was collected with written, prior, informed consent and the data collection was reviewed by the university ethics board.

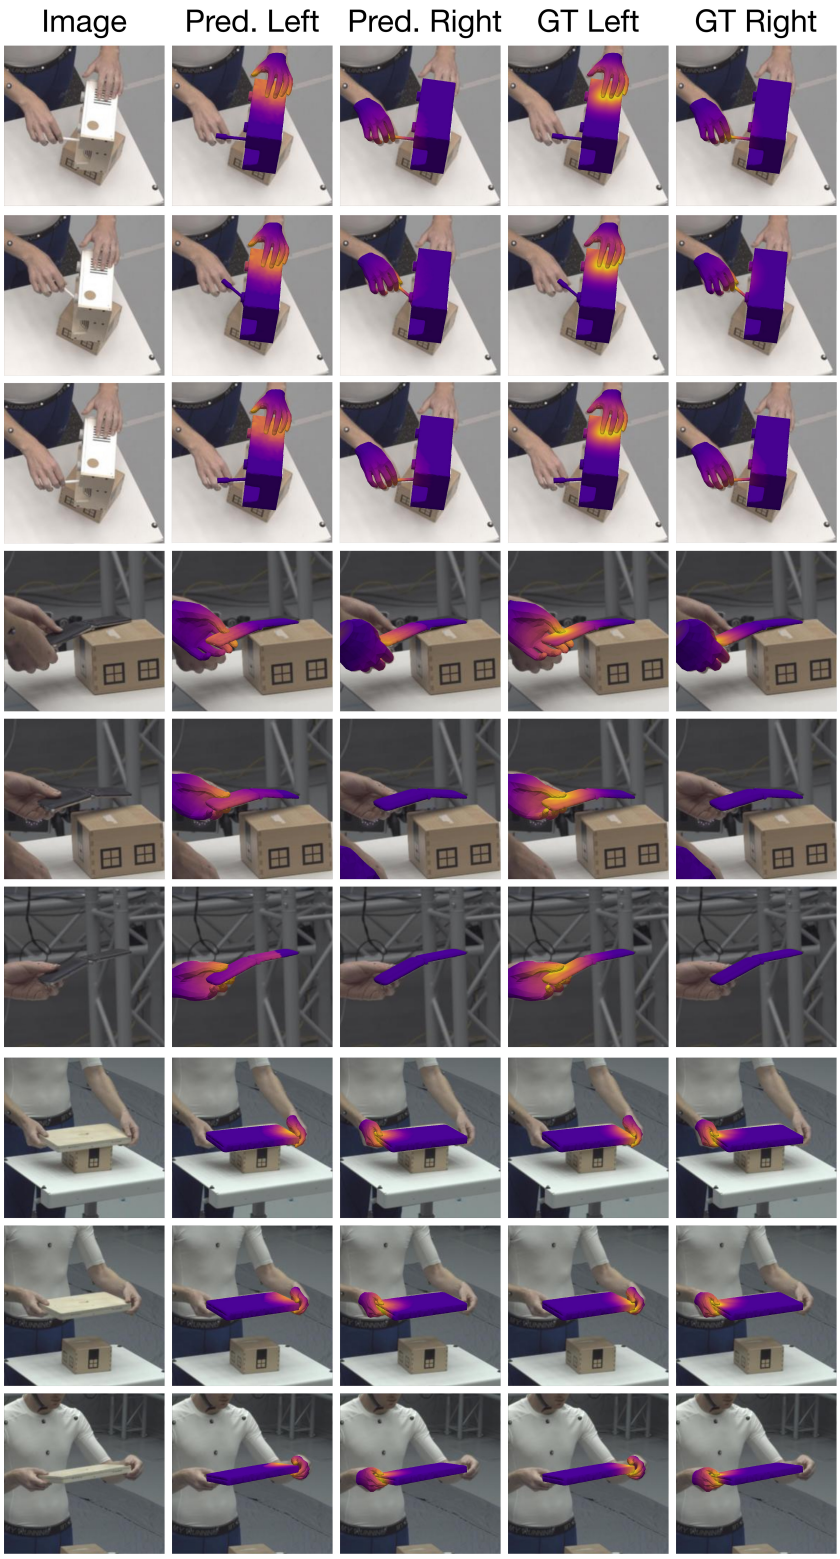


Figure 9. **Qualitative results of InterField-SF**. Best viewed in color and zoomed in.

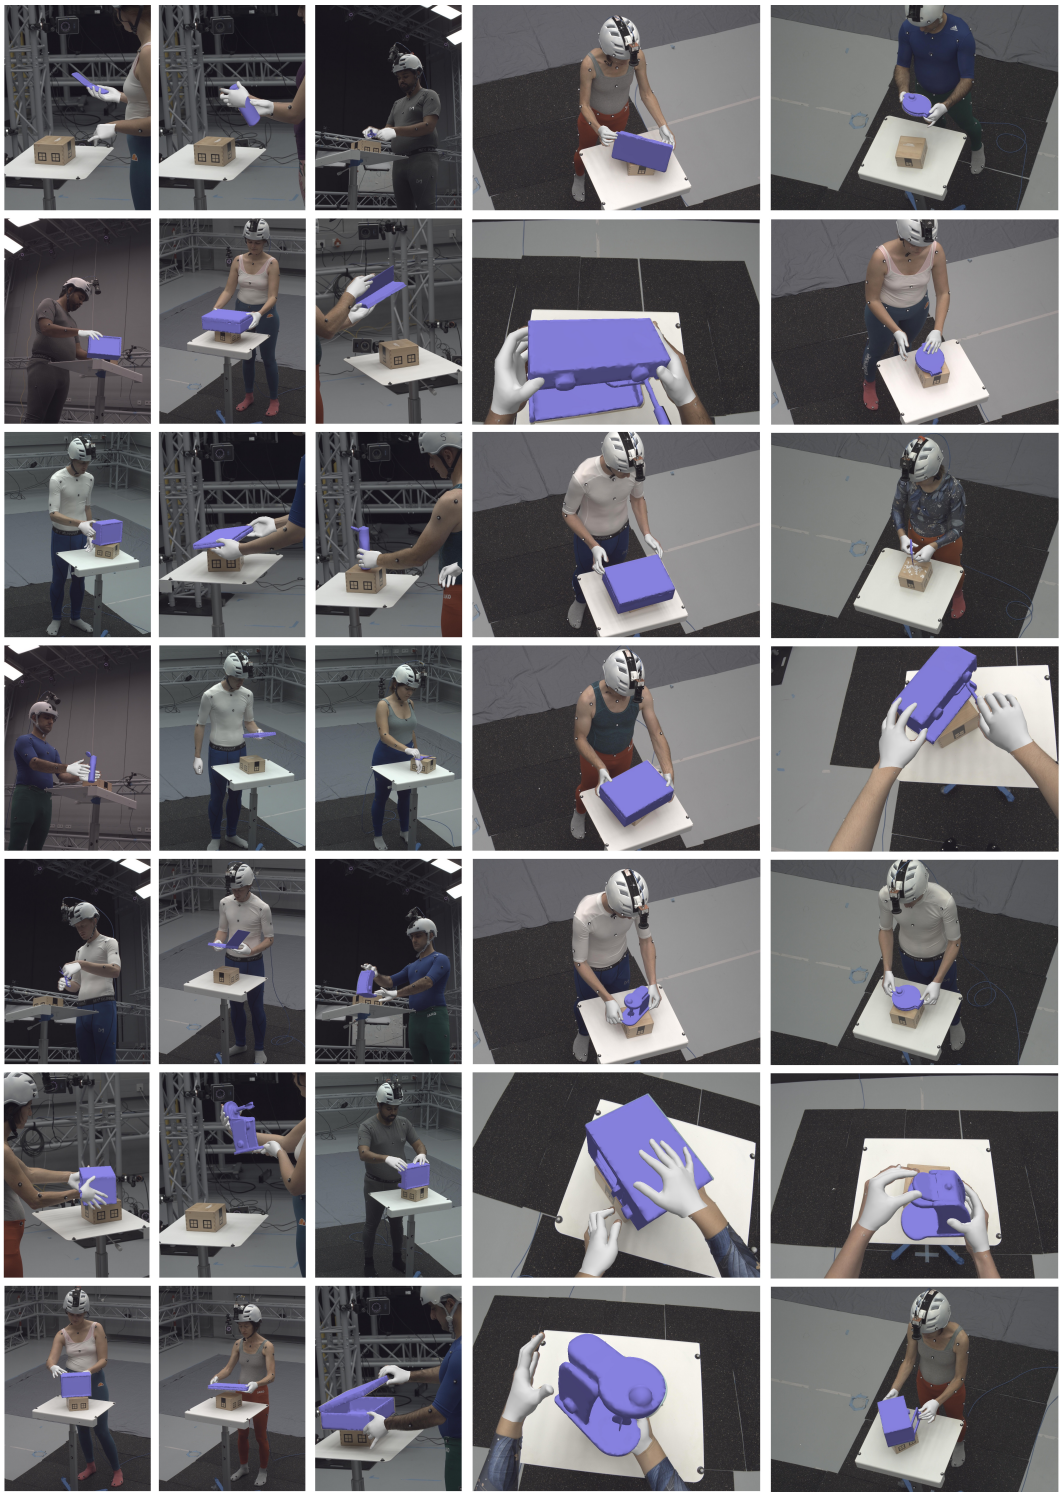


Figure 10. **Overlay of ARCTIC ground-truth** We overlay the ground-truth in our dataset. Examples here are randomly sampled from ARCTIC. Note that although the RGB images contain both human bodies and hands, the hand region is clearly visible when zoomed in, thanks to our high resolution RGB images.