

PokeFlex: A Real-World Dataset of Volumetric Deformable Objects for Robotics

Jan Obrist*, Miguel Zamora*, Hehui Zheng*, Ronan Hinchet,
Firat Ozdemir, Juan Zarate, Robert K. Katzschmann, and Stelian Coros

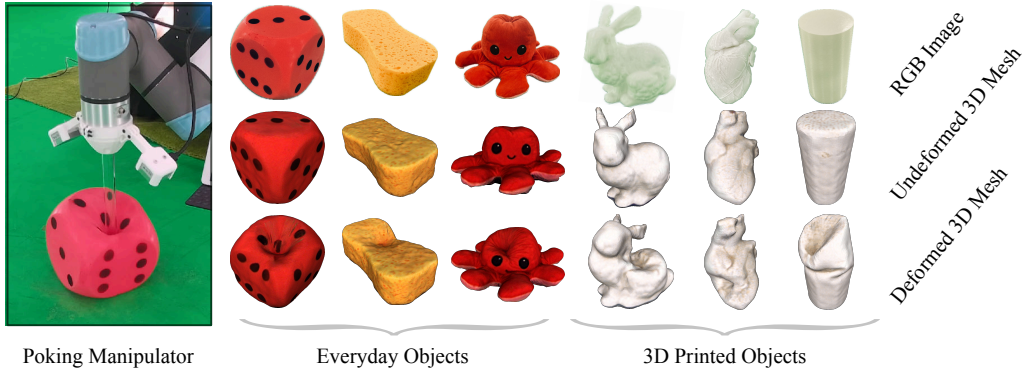


Fig. 1: PokeFlex captures the deformability of various everyday and 3D-printed objects, as illustrated by the poking manipulator on the **Left**. On the **Right**, the **Top Row** contains segmented RGB images of selected objects. The **Middle Row** shows reconstructed objects in an undeformed state. The **Bottom Row** provides reconstructed 3D-textured meshes of deformed objects. For an overview of all the objects, see Figure 6.

Abstract—Data-driven methods have shown great potential in solving challenging manipulation tasks; however, their application in the domain of deformable objects has been constrained, in part, by the lack of data. To address this lack, we propose PokeFlex, a dataset featuring real-world multimodal data that is paired and annotated. The modalities include 3D textured meshes, point clouds, RGB images, and depth maps. Such data can be leveraged for several downstream tasks, such as online 3D mesh reconstruction, and it can potentially enable underexplored applications such as the real-world deployment of traditional control methods based on mesh simulations. To deal with the challenges posed by real-world 3D mesh reconstruction, we leverage a professional volumetric capture system that allows complete 360° reconstruction. PokeFlex consists of 18 deformable objects with varying stiffness and shapes. Deformations are generated by dropping objects onto a flat surface or by poking the objects with a robot arm. Interaction wrenches and contact locations are also reported for the latter case. Using different data modalities, we demonstrated a use case for our dataset training models that, given the novelty of the multimodal nature of Pokeflex, constitute the state-of-the-art in multi-object online template-based mesh reconstruction from multimodal data, to the best of our knowledge. The full dataset, pretrained models, and code are available on the project website¹.

Index Terms—Data Sets for Robotic Vision; Deep Learning for Visual Perception; Data Sets for Robot Learning

I. INTRODUCTION

THE development of high-quality datasets is essential to advance research in deformable object manipulation using data-driven methods, which have recently demonstrated promising results in fields such as healthcare, food processing, and manufacturing [1]–[4]. Such datasets are crucial for training manipulation policies, estimating material parameters, and training 3D mesh reconstruction models. The latter, in particular, plays a vital role in facilitating the close-loop execution of control methods based on mesh simulations [5]. In light of these needs, the objective of this work is to create a reproducible, diverse, and high-quality dataset for deformable volumetric objects, grounded in real-world data.

Current state-of-the-art simulation methods offer an attractive alternative to collect such datasets, providing easy access to privileged information such as deformed mesh configurations and contact forces [10]–[15]. However, such simulators require careful system identification and fine-tuning to bridge the sim-to-real gap, which ultimately requires real-world data. Static scans rotating around the scene [9], [16] or custom multi-camera systems [7] can be used to collect real-world 3D models. The former can be excessively time-consuming and is unsuitable to capture temporal dynamics. The latter requires careful synchronization and data curation, especially when using noisy lower-cost sensors.

To address these challenges, we leverage a professional multi-view volumetric capture system (MVS) that allows capturing detailed 360° mesh reconstructions of deformable objects over time [17], which we use as ground-truth meshes.

Manuscript received: January, 23, 2025; Revised May, 13, 2025; Accepted July, 19, 2025. This paper was recommended for publication by Editor A. Faust upon evaluation of the Associate Editor and Reviewers' comments. This work is supported by the SDSC Grant entitled 'C22-08: Data-Driven Inference of Mesh-based Representations for Deformable Objects from Unstructured Point Clouds'. * Denotes equal contribution.

J. Obrist*, H. Zheng*, R. Hinchet, and R. K. Katzschmann are with the Department of Mechanical and Process Engineering, ETH Zurich, Switzerland rkk@ethz.ch

M. Zamora*, J. Zarate, and S. Coros are with the Department of Computer Science, ETH Zurich, Switzerland scoros@ethz.ch

F. Ozdemir is with the Swiss Data Science Center, ETH Zurich & EPFL, Zurich, Switzerland.

Digital Object Identifier (DOI): see top of this page.

¹<https://pokeflex-dataset.github.io/>

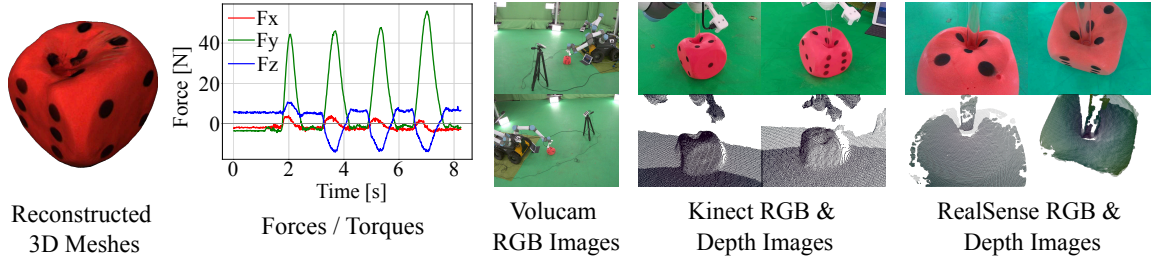


Fig. 2: Different data modalities provided by PokeFlex.

TABLE I: Dataset overview (per object, per sequence).

Sequence Data	Poking	Dropping
• 3D textured deformed mesh model	✓	✓
• RGB images from two Volucam cameras (cameras from the MVS)	✓	✓
• RGB-D images from two Intel RealSense D405 sensors (eye-in-hand mounted)	✓	
• RGB-D images from two Azure Kinect sensors (eye-to-hand mounted)	✓	
• Estimated 3D contact forces and torques	✓	
• End-effector poses and contact locations	✓	
Camera and Object Data		
• Camera intrinsic and extrinsic parameters		✓
• 3D textured template mesh model		✓
• Open-source print files to reproduce the 3D printed objects		✓

We integrate a robotic manipulator with joint-torque sensing capabilities into the MVS, enabling contact force estimation and automated data collection. Moreover, to enhance reproducibility and to expand the diversity of data modalities, we also integrate and synchronize lower-cost Azure Kinect and Intel RealSense D405 RGB-D sensors into the MVS.

Our work proposes the PokeFlex dataset (Figure 1), featuring the real-world behavior of 18 deformable objects, including everyday and 3D-printed objects. Deformations are generated via controlled poking and dropping protocols. An overview of the paired, synchronized, and annotated data is illustrated in Figure 2, and summarized in Table I. We demonstrated a use case of the PokeFlex dataset, proposing baseline models capable of ingesting PokeFlex multimodal data and presenting evaluation criteria for benchmarking the results. Specifically, we train neural network models for deformed mesh reconstructions based on template meshes and various input data modalities, including images, point clouds, end-effector poses and forces. The proposed architectures are suitable for online applications, reconstructing 3D meshes at a range from 33 Hz to 185 Hz depending on the input data modality, on a desktop PC with an NVIDIA RTX 4090 GPU. The pretrained models are available on PokeFlex website.

II. RELATED WORK

Deformable object datasets. Depending on the use of synthetic or real-world data, deformable object datasets can be roughly categorized into two major groups. [10], for instance, evaluates multiple grasping poses for deformable

objects on a large-scale synthetic dataset. Qualitative sim-to-real experiments for such dataset, show that their simulator captures the general deformation behavior of objects during grasping. However, careful system identification and parameter tuning are necessary to achieve higher sim-to-real fidelity for synthetic datasets.

On the other hand, real-world data collection opens up the door to better capture the complex behavior of deformable objects. Current real-world datasets focus mostly on RGB images. HMDO [6] also provides real-world 3D meshes for objects undergoing deformation due to hand manipulation. However, they fell short of providing point cloud or force contact information. [7] provides points clouds and force contact information but it does not perform 3D mesh reconstruction and the deformations are only globally produced using an airstream. [8] offer a large number of frames, however, the object diversity in their dataset is limited.

In a departure from other datasets, PokeFlex offers a more comprehensive list of features including; 3D meshes, point clouds, contact forces, higher diversity of objects, and multiple types of deformations as detailed in Table II. For simplicity, we report only the effective number of paired time frames in our table, in contrast to what is reported by [6] and [8], where the total number of samples is computed as the number of time-frames times the number of cameras.

Data-driven mesh reconstruction methods are particularly useful for manipulating deformable objects, as they provide a means to estimate the partially observable internal state of such objects. In the domain of cloth manipulation, prior work [18]–[20] has demonstrated how mesh reconstruction facilitates manipulation tasks. Even in the case of rigid objects, 3D mesh reconstruction has been used as an informative prior to train manipulation policies [21] and to improve the performance of model-based controllers [22]. Reconstruction methods vary widely in terms of the input data modalities they employ. Previous approaches that rely on point clouds to predict deformations are typically trained on synthetic data [23]–[25]. While synthetic training data offers controlled and dense point cloud representations, it often leads to a sim-to-real gap as real-world point cloud measurements tend to be noisy and sparse, especially in dynamic and unstructured environments. In contrast, methods using single images as input have gained attention for their real-world reconstruction capability without the need for depth information [26]–[28]. However, many of these image-based approaches are not optimized for online inference, making them unsuitable for

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE II: Feature comparison of the PokeFlex dataset with other deformable object datasets.

	Real-world	Meshes	Point clouds	RGB images	Force torque	# of objects	# of time frames	Type of deformation
PokeFlex (ours)	✓	✓	✓	✓	✓	18	21.3k	Poke, drop
HMDO [6]	✓	✓		✓		12	2,166	Manual [†]
PLUSH [7]	✓		✓	✓	Force [‡]	12	22.84k	Airstream
DOT [8]	✓		✓	✓		4	117k	Manual
Household Cloth Object Set [9]	✓	✓ [§]		✓		27	67	/
Defgraspsim [10]		✓				34	1.1M	Grasp

[†] by hand

[‡] by providing air nozzle poses

[§] for ten static scenes of the cloth objects folded




Fig. 3: Sample frames from a poking sequence, with a close-up onto the foam dice.

downstream applications in robotics, where online feedback is essential. An instant image-to-3D framework to generate high-quality 3D assets is proposed by [29], but it requires up to 10 seconds per frame, limiting its practicality for real-time processing scenarios.

III. METHODOLOGY

A. Data Acquisition

The PokeFlex dataset involves the acquisition of deformations under two different protocols (i) poking and (ii) dropping. In the poking protocol, a robotic manipulator pokes objects with a transparent acrylic stick multiple times along a randomly oriented horizontal vector (Figure 3). The dataset also provides the CAD model for the mounting tool, which holds two RealSense cameras and a 192 mm long acrylic stick with a radius of 10 mm. In the dropping protocol, objects are attached to a light nylon cord at approximately 2 m height and captured in a free-fall drop onto a flat surface. We record data at 30 fps and 60 fps for the poking and dropping protocols, respectively. We leverage a professional multi-view volumetric capture system (MVS), consisting of 106 cameras (53 RGB / 53 infrared) with 12 MP resolution.

For the poking protocol, we integrated and synchronized additional hardware to the MVS to ensure temporally aligned data capture across all modalities. The additional hardware includes the robot manipulator and four additional RGB-D cameras: two Azure Kinect cameras to capture the scene from opposing viewpoints, and two Intel RealSense D405 cameras mounted on the robot’s end-effector. The robot logs end-effector poses, interaction forces and torques at 120 Hz, while these four cameras record RGB-D data at 30 Hz.

To synchronize devices, we rely on a Linear Timecode (LTC) signal provided by an Atomos Ultrasync device. The cameras of the MVS have a leader/follower architecture, where the internal clocks of the follower cameras are synchronized to one single leader camera, which reads the LTC signal. In addition to the MVS control system, we use two desktop PCs

to read the additional data streams: a Robot PC that reads the robot data and the streams of the two RealSense D405 cameras and a dedicated Kinect PC that reads the streams of the two Azure Kinect devices.

The Robot PC is synchronized with the capture system by reading the same LTC signal provided by the Atomos Ultrasync device. The Kinect cameras are hardware-synchronized with each other. Their synchronization with the capture system is achieved retrospectively by comparing the current timecode displayed on a screen in the camera frames of the Kinect and the camera frames of the capture system. An overview of the architecture is shown on Figure 4.

We utilize a system similar to that described by [17] to reconstruct the meshes and textures of the objects under deformation. When recording at 30 fps, the MVS generates approximately 27 GB of raw data per second. This data is then processed using commercial software provided by Acturus Studio on 10x On-Prem Nodes servers, achieving an output rate of approximately one 3D frame per minute. The reconstructed meshes and textures were curated to ensure that only the deformable objects were retained in the scene.




Fig. 4: **Top:** Robotic manipulator positioned inside MVS with external lower-cost cameras sensors. **Bottom:** Overview of the system architecture to capture PokeFlex data.

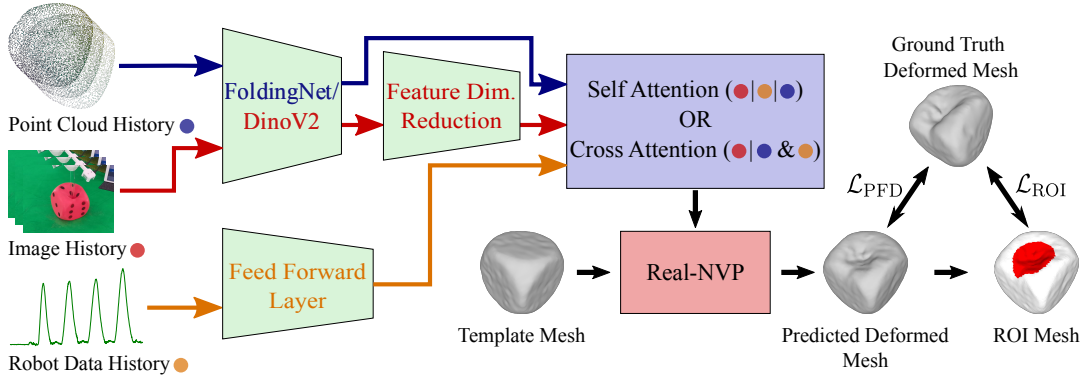


Fig. 5: Superimposed representation of the proposed network architectures for ingesting the multi-modal PokeFlex data.

B. Learning-based Mesh Reconstruction

We leverage PokeFlex to train models for template-based mesh reconstruction, inferring the deformation of a template mesh from an input of history buffers with various combinations of data modalities: images, point clouds, and/or robot data. Figure 5 illustrates the building blocks used to generate different architectures tailored to the data modalities.

At a high level, we use three main common components for all models: an encoder to extract features from an input modality, an attention mechanism to exploit temporal information from the history buffers, and a conditional Real-NVP [23] to predict the offsets of template vertices, yielding the predicted deformed mesh. Real-NVP uses a series of conditional coupling blocks, defined as continuous bijective functions. This continuous bijective operation ensures that the model is homeomorphic, allowing stable deformation of a template mesh M_T , while preserving its topology.

Image input: For pipelines using images as input, we use a DinoV2 vision transformer to extract embeddings of each image frame. In particular, we use a DinoV2-small model, pretrained via distillation from the largest DinoV2 transformer presented in [30] (LVD-142M dataset). The embedding dimension is later reduced using a 1D convolutional layer and a subsequent fully connected layer (Feature Dim. Reduction block in Figure 5).

Point cloud input: When using point clouds, we leverage a FoldingNet encoder [31] for representation learning, which is trained end-to-end together with the attention mechanism and the conditional-NVP.

Robot data input: To fuse the robot data, we concatenate the measured end-effector forces and the position of the interaction point. The concatenated data is later fed into a single fully connected layer, to match the dimensionality of the embeddings used for the attention mechanisms.

A self-attention mechanism is employed for variations of the architecture in Figure 5 that use a single data modality as input. In contrast, a cross-attention mechanism is applied when handling multiple data modalities simultaneously. For the experiments presented in the results section, we use cross-attention to handle a mixture of history buffers of robot data with images or Kinect point clouds, as input. However, other combinations of input data are also possible.

All architectures are end-to-end trained using the same loss. We include the weights of the DinoV2 transformer during backpropagation for finetuning. The main point face distance (PFD) criterion \mathcal{L}_{PFD} accounts for the global deformation of the objects, which computes the average squared distance $d(\mathbf{p}, \mathbf{f})$ from the set of sampled points $\mathbf{p}_i \in \mathcal{P}$ on the predicted mesh to the nearest faces in the set of triangular faces $\mathbf{f}_j \in \mathcal{F}$ of the ground truth mesh and vice versa (eq. (1)). Moreover, to deal with the local deformations generated in the poking region, we add a region-of-interest (ROI) loss \mathcal{L}_{ROI} (eq. (2)) that computes the unidirectional chamfer distance from the points \mathbf{p}_i in the ROI to the set of sampled points $\mathbf{q}_j \in \mathcal{Q}$ of the ground truth mesh. The ROI is defined using the indicator function $\mathbb{I}(\mathcal{C}(\mathbf{p}_i))$, which evaluates to 1 if point \mathbf{p}_i is close enough to the contact point \mathbf{t} according to a threshold ϵ , and if the minimum vertical component of the contact point $\mathbf{p}_{i,y}$ is bigger than a vertical threshold ϵ_y (eq. (3)).

$$\mathcal{L}_{\text{PFD}} = \frac{1}{|\mathcal{P}|} \sum_{\mathbf{p}_i \in \mathcal{P}} \min_{\mathbf{f}_j \in \mathcal{F}} d(\mathbf{p}_i, \mathbf{f}_j) + \frac{1}{|\mathcal{F}|} \sum_{\mathbf{f}_j \in \mathcal{F}} \min_{\mathbf{p}_i \in \mathcal{P}} d(\mathbf{f}_j, \mathbf{p}_i), \quad (1)$$

$$\mathcal{L}_{\text{ROI}} = \frac{1}{|\mathcal{P}|} \sum_{\mathbf{p}_i \in \mathcal{P}} \mathbb{I}(\mathcal{C}(\mathbf{p}_i)) \cdot \min_{\mathbf{q}_j \in \mathcal{Q}} \|\mathbf{p}_i - \mathbf{q}_j\|^2, \quad (2)$$

$$\mathcal{C}(\mathbf{p}_i) = (\|\mathbf{p}_i - \mathbf{t}\| \leq \epsilon) \wedge (\mathbf{p}_{i,y} > \epsilon_y). \quad (3)$$

The total loss is then set as $\mathcal{L} = \mathcal{L}_{\text{PFD}} + 0.5 \mathcal{L}_{\text{ROI}}$.

IV. RESULTS

A. Dataset

The PokeFlex dataset comprises 18 deformable objects (Figure 6), including 13 everyday items as well as 5 objects that are 3D printed with a soft thermoplastic polyurethane filament. Even though the everyday objects in our dataset can be purchased from global vendors, their availability is not guaranteed worldwide. Therefore, to enhance the usability of our dataset we include deformable 3D printed objects, providing print files and detailed specifications for reproducibility (reported on our website). The 3D printed objects include the Stanford bunny [32], a cylinder, a heart [33], a pyramid, and a custom pizza slice.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

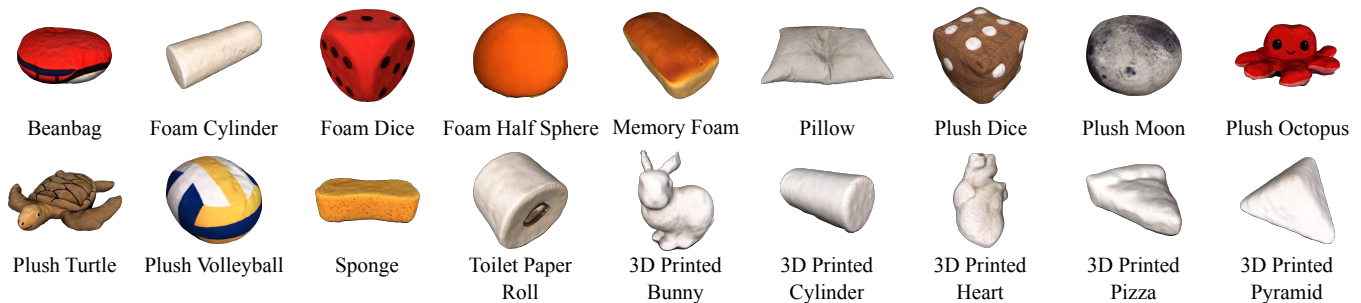


Fig. 6: Rest-state reconstructed 3D meshes of all 18 objects featured in the PokeFlex dataset.

The dimensions and the weights of the PokeFlex objects range from 7 cm to 58 cm and from 22 g to 1 kg, respectively. Furthermore, using Hooke’s law and applying RANSAC for linear regression to avoid outliers, we estimated the objects’ stiffnesses to be in the range of 148–2,156 N/m.

For the poking protocol, we recorded 4-8 sequences with a duration of 5-6 seconds at 30 fps for each object. Similarly, for the dropping protocol, we recorded 3 sequences of 1 second at 60 fps for each object. Figure 7 shows two reconstructed sequences for poking and dropping. In the case of the poking sequences, each frame includes synchronized and paired data from all modalities, as illustrated in Figure 2.

The total number of reconstructed frames used to generate ground-truth data was 21.3k, which comprises 18.1k frames for the poking sequences and 3.2k frames for the dropping sequences. It is worth noting that after curating the frames of the poking sequences, i.e., discarding the frames where the robot arm is not in contact with the objects, the total number of active paired poking frames reduces from 18.1k down to 8.6k. Furthermore, if we consider all the different modalities to count the number of samples as done in [6], [8], the total number of samples in PokeFlex amounts to 240k. A summary of the physical properties of the objects, as well as a per-object list of the recorded frames under deformation for the poking sequences, is presented in our website. For the dropping protocol, we recorded 180 frames per object.

B. Evaluation of Learning-based Reconstruction

Overview of training data. In the following experiments, we exclusively used poking sequences from the dataset be-




Fig. 7: **Top:** Mesh reconstructions of foam dice for a poking sequence shown in every third frame. **Bottom:** Mesh reconstructions of plush octopus for a dropping sequence.

cause of the higher diversity of input data modalities available. The size of the history buffer was ablated and set to 5 for better performance. The train-validation split was generated by randomly choosing one recording sequence per object as the validation set. Additional training hyper-parameters are reported in Appendix A.

Metrics. During training, we reposition and re-scale all meshes into a cube of unit size $([-0.5, 0.5]^3)$ to maintain a consistent scale across all objects. The losses \mathcal{L}_{PFD} and \mathcal{L}_{ROI} are computed in this normalized scale. To further assess the prediction accuracy, we evaluate two additional metrics between the predicted mesh M_P and the ground truth mesh M_{GT} in their original scale: the unidirectional L1 Norm Chamfer Distance CD_{ULI} (eq. (4)) and the volumetric Jaccard Index J (eq. (5)), which we defined in terms of the volume operator V . The metrics provide insights into the L1 Norm surface distance and the volume overlap ratio, respectively.

$$\text{CD}_{\text{ULI}} = \frac{1}{|\mathcal{P}|} \sum_{\mathbf{p}_i \in \mathcal{P}} \min_{\mathbf{q}_j \in \mathcal{Q}} \|\mathbf{p}_i - \mathbf{q}_j\|_1, \quad (4)$$

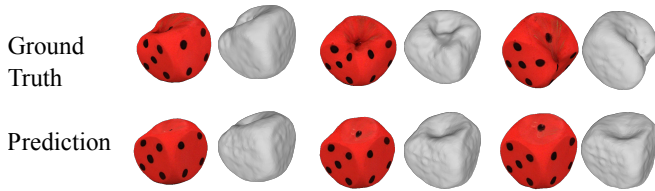
$$J(M_A, M_B) = \frac{V(M_A \cap M_B)}{V(M_A \cup M_B)}. \quad (5)$$

Learning from different data modalities. In this experiment, we train different mesh prediction models from sequences of different input modalities. We trained multi-object models using all 18 objects from the dataset. Detailed

TABLE III: Mean prediction performance for proposed model configurations trained on all objects. Arrows indicate that a better performance is either higher \uparrow or lower \downarrow . ($J := J(M_P, M_{\text{GT}})$, $\text{PCD} := \text{Point Clouds}$)

Input	$\mathcal{L}_{\text{PFD}} \downarrow$ $\cdot 10^3$	$\mathcal{L}_{\text{ROI}} \downarrow$ $\cdot 10^3$	$\text{CD}_{\text{ULI}} \downarrow$ [mm]	$J \uparrow$	Speed \uparrow [Hz]
Ours					
• Images	7.34	7.22	7.547	0.791	115
• Robot data	8.47	5.50	8.283	0.779	185
• Images + robot data	6.64	4.99	7.025	0.806	110
• Dense ground truth PCD	4.34	3.64	6.185	0.825	33
• Sparse ground truth PCD [¶]	4.84	4.77	6.445	0.823	180
• Kinect PCD ^{**}	5.75	5.00	6.498	0.820	33
• Kinect PCD + robot data	5.68	4.06	6.615	0.817	33
SAP (optimization-based) [34] - 100 iterations					
• Dense ground truth PCD	14.0	10.6	9.641	0.749	0.20
• Sparse ground truth PCD	116	32.5	13.78	0.238	0.20
• Kinect PCD	62.3	18.6	29.51	0.377	0.19
5k points [¶] 100 points ^{**} 5k points					

performance for the evaluated data modalities can be found in Table III, as well as their corresponding inference rates, which range from 33 Hz to 185 Hz for dense point clouds and robot data, respectively. Figure 8 shows examples of predicted meshes with different levels of reconstruction quality obtained using a multi-object model trained from image-sequences only. Ground truth point clouds (PCDs) were generated by sampling the meshes reconstructed from real-world data with the MVS, which we considered as ground truth meshes. As such, the results obtained from ground truth PCDs are meant to illustrate the performance that could be achieved in an ideal setting from noise-free PCDs. Furthermore, Table III also includes optimization-based Shape-As-Points (SAP) [34], which is a Poisson Surface Reconstruction method, to highlight the benefits of our approach in terms of reconstruction quality and speed.



Performance	Good	Medium	Outlier
$\mathcal{L}_{\text{PFD}} * 10^3$	4.21	5.91	12.35
$\mathcal{L}_{\text{ROI}} * 10^3$	5.17	5.00	6.60
$\text{CD}_{\text{ULI}} [\text{mm}]$	5.63	6.55	9.64
$J(M_P, M_{GT})$	0.87	0.86	0.79

Fig. 8: Examples of deformation predictions for a foam dice and their corresponding metrics. Meshes are rendered side by side with and without texture to highlight the ROI.

Generalizing to unseen objects. To test the generalization capabilities of the proposed architectures, we trained a model on ground truth point cloud data (Table IV) on 14 different objects and evaluated on 4 unseen objects. The generalization performance was evaluated using all the available sequences from the unseen objects.

TABLE IV: Generalization results for 4 unseen objects for point-cloud-based mesh reconstruction.

Input	$\mathcal{L}_{\text{PFD}} \downarrow$ ·10 ³	$\mathcal{L}_{\text{ROI}} \downarrow$ ·10 ³	$\text{CD}_{\text{ULI}} \downarrow$ [mm]	$J(M_P, M_{GT}) \uparrow$
• Validation set (14 objects)	5.13	4.26	6.247	0.816
• Foam cylinder	4.77	2.42	8.237	0.806
• Plush dice	2.94	3.59	7.838	0.895
• Plush volleyball	2.48	3.52	6.074	0.903
• Sponge	9.70	3.54	8.673	0.711

V. DISCUSSION

Quality of ground-truth meshes. The overall geometry of the objects in the dataset, in static configurations, is well captured by the meshes reconstructed with the MVS as shown in Figure 6, even though the system’s intended use is the reconstruction of human-size objects. Furthermore, the proposed poking protocol, using a transparent acrylic



Fig. 9: Examples of reconstructed ground truth meshes for medium (Left) and small (Right) size objects in deformed states. Reconstruction of fine-grained details is a limitation of our current setup (close-up views on the Right).

stick, helps prevent occlusions at the contact point, leading to detailed reconstruction of objects even when they undergo deformations, as can be seen in Figure 9 (Left). However, reconstruction of fine-grained details for smaller objects such as the 3D-printed Stanford armadillo [35] remains challenging with the current setup of the MVS (Figure 9 - Right). Better fine-grained reconstruction results can be expected by rearranging the cameras in a smaller workspace, by leveraging high-resolution templates, or physics-based priors for 3D mesh reconstruction [36].

Multi-object mesh reconstruction from different modalities. Table III shows that the dense ground truth point clouds yield the best performance among all data modalities. A drop in performance is observed for the sparser ground truth point clouds, and the noisier point clouds captured by the Kinect. The models trained from robot data and images or Kinect point clouds generally outperformed the counterparts trained from only robot data, only images, or only Kinect point clouds, showcasing the importance of combining different input modalities, and hinting at the effectiveness of our cross-attention mechanism. In addition, we can observe that the CD_{ULI} metric correlates positively with the training losses and it serves as a more intuitive way to evaluate performance in the original scales of the deformable objects. Across all evaluation metrics, our method outperformed the optimization-based Shape-As-Points (SAP) approach [34]. While SAP can produce better reconstruction results with increased iterations, the associated computational cost makes SAP even more unsuitable for online applications.

To analyze the levels of accuracy across multiple objects, we focus on the image-based mesh reconstruction model. Figure 10 illustrates the performance of our model in terms of the Jaccard index $J(M_P, M_{GT})$ vs the Jaccard distance, which is defined as $d_J(M_T, M_{GT}) = 1 - J(M_T, M_{GT})$ and indicates the level of deformation of the ground truth mesh with respect to the template mesh. A perfect model would predict meshes that result in a horizontal line at $J = 1.0$. For our image-based mesh reconstruction model, we note that the Jaccard Index presents, in general, a slight negative correlation with the Jaccard distance for all objects, indicating that the prediction accuracy decreases for larger deformations. The results in Figure 10 also offer insights into the per-object performance, where the best results are obtained for objects like the foam dice and the foam half sphere, while objects like the plush turtle and the memory foam present a more

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

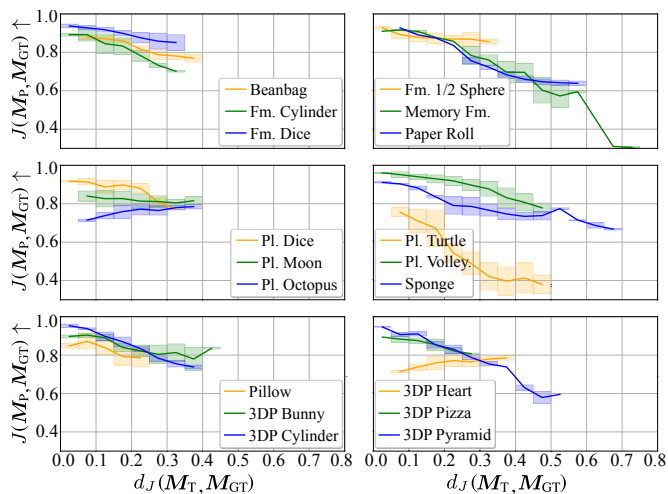


Fig. 10: Validation accuracy for image-based mesh reconstruction, evaluated by Jaccard Index J , plotted against the deformation level quantified by Jaccard distance d_J . (Fm.::= Foam, Pl.::= Plush, 3DP:= 3D Printed, Volley.::= Volleyball)

significant challenge for our image-based mesh reconstruction model.

Reconstruction performance on unseen objects To evaluate the generalization of the ground truth point cloud mesh reconstruction model, we use the mean validation performance achieved on the 14 objects that belonged to the training set as a baseline. In terms of the point face distance loss \mathcal{L}_{PFD} , the model demonstrates a good level of generalization for objects like the foam cylinder, the plush dice, and the plush volleyball, compared to the baseline. However, in terms of the unidirectional chamfer distance CD_{UL1} , only the plush volleyball outperforms the baseline. As expected, performance tends to degrade on unseen objects. Nonetheless, the plush volleyball shows good generalization across all metrics, which could be attributed to the fact that it has a similar shape and size to other objects in the training set.

VI. FUTURE WORK

The 3D meshes captured during poking and dropping are well-suited for system identification of material parameters in physics-based simulation models [37]. These can be leveraged using sampling-based methods for non-differentiable simulators like Isaac Sim [38], or through gradient-based optimization in differentiable simulators [39], [40]. Alternatively, the PokeFlex dataset can also be used to learn dynamic models [41], [42] directly from real-world data. Developing fast, stable, and realistic simulation models is key to enabling the deployment of learning-based or model-based controllers in real-world applications, either by training reinforcement learning (RL) policies in simulation or by employing model predictive control (MPC) strategies. The latter would require perception models capable of estimating the full mesh-based state of volumetric deformable objects. The colored textures of our 3D meshes can be used to render synthetic images from multiple static or dynamic viewpoints, providing valuable data for training view-agnostic perception models. Applying these

textures to deformable objects in simulated pick-and-place or non-prehensile manipulation tasks also enables scalable synthetic visual data generation within physics-based environments.

Future iterations of the dataset could extend to dual-arm pick-and-place tasks, using two transparent acrylic sticks, similar to our current setup, or by employing transparent 3D-printing technology [43] for parallel grippers or open-source dexterous hands [44]. Ultimately, achieving general-purpose dexterous manipulation of deformable objects, beyond traditional parallel grips and power grasps, will likely depend on new algorithms that can learn from a combination of real-world data, physics-based simulations, and synthetic data generated from world foundation models.

VII. CONCLUSION

This paper introduced PokeFlex, a new dataset that captures the behavior of 18 deformable volumetric objects during poking and dropping. The focus is on volumetric objects, while thin clothing items or thin cables are not considered in the dataset. Compared to previously existing datasets, we provide a wider range of paired and annotated data modalities, which are supplemented with data streams from lower-cost camera sensors. In an effort to enhance reproducibility, the objects included in our dataset can be either purchased from global providers or 3D printed with our open-source models.

Using different combinations of the data modalities provided in PokeFlex, we train and benchmark a list of baseline models, which constitutes the state-of-the-art for the task of multi-object online template-based mesh reconstruction, given the novelty of the multimodal nature of our dataset, to the best of our knowledge. We are excited to share this dataset with the community and hope it serves as a valuable resource to push the boundaries of robotics research.

APPENDIX

A. Training Details

TABLE V: Training hyperparameters.

Hyperparameters	Value
Learning rate	1e-4
Batch size	32
Optimizer	Adam
Weight decay	0.0
Learning rate scheduler	Cosine
Minimum learning rate	1e-7
Epochs	200

REFERENCES

- [1] Q. Yu, M. Moghani, K. Dharmarajan, V. Schorp, W. C.-H. Panitch, J. Liu, K. Hari, H. Huang, M. Mittal, K. Goldberg, and A. Garg, “Orbitsurgical: An open-simulation framework for learning surgical augmented dexterity,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 15 509–15 516.
- [2] H. Shi, H. Xu, S. Clarke, Y. Li, and J. Wu, “Robocook: Long-horizon elasto-plastic object manipulation with diverse tools,” *arXiv preprint arXiv:2306.14447*, 2023.
- [3] Y. Avigal, L. Berscheid, T. Asfour, T. Kröger, and K. Goldberg, “Speedfolding: Learning efficient bimanual folding of garments,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 1–8.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

- [4] A. Bartsch, C. Avra, and A. B. Farimani, "Sculptbot: Pre-trained models for 3d deformable object manipulation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 12 548–12 555.
- [5] S. Duenser, J. M. Bern, R. Poranne, and S. Coros, "Interactive robotic manipulation of elastic objects," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 3476–3481.
- [6] W. Xie, Z. Yu, Z. Zhao, B. Zuo, and Y. Wang, "Hmdo : Markerless multi-view hand manipulation capture with deformable objects," *Graphical Models*, vol. 127, p. 101178, 2023.
- [7] H.-y. Chen, E. Tretschk, T. Stuyck, P. Kadlecck, L. Kavan, E. Vouga, and C. Lassner, "Virtual elastic objects," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 15 827–15 837.
- [8] X. Li, Y. Guo, Y. Tu, Y. Ji, Y. Liu, J. Ye, and C. Zheng, "Textureless deformable object tracking with invisible markers," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [9] I. Garcia-Camacho, J. Borràs, B. Calli, A. Norton, and G. Alenyà, "Household cloth object set: Fostering benchmarking in deformable object manipulation," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 5866–5873, 2022.
- [10] I. Huang, Y. Narang, C. Eppner, B. Sundaralingam, M. Macklin, R. Bajcsy, T. Hermans, and D. Fox, "Defgraspsim: Physics-based simulation of grasp outcomes for 3d deformable objects," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6274–6281, 2022.
- [11] P. Tripicchio, S. D'Avella, and C. A. Avizzano, "Cepb dataset: a photorealistic dataset to foster the research on bin picking in cluttered environments," *Frontiers in Robotics and AI*, vol. 11, p. 1222465, 2024.
- [12] M. Macklin, "Warp: A high-performance python framework for gpu simulation and graphics," <https://github.com/nvidia/warp>, March 2022, nVIDIA GPU Technology Conference (GTC).
- [13] Y. Qiao, J. Liang, V. Koltun, and M. Lin, "Differentiable simulation of soft multi-body systems," *Advances in Neural Information Processing Systems*, vol. 34, pp. 17 123–17 135, 2021.
- [14] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 5026–5033.
- [15] F. Faure, C. Duriez, H. Delingette, J. Allard, B. Gilles, S. Marchesseau, H. Talbot, H. Courtecuisse, G. Bousquet, I. Peterlik *et al.*, "Sofa: A multi-model framework for interactive physical simulation," *Soft tissue biomechanical modeling for computer assisted surgery*, pp. 283–321, 2012.
- [16] D. K. Pai, K. v. d. Doel, D. L. James, J. Lang, J. E. Lloyd, J. L. Richmond, and S. H. Yau, "Scanning physical interaction behavior of 3d objects," in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 2001, pp. 87–96.
- [17] A. Collet, M. Chuang, P. Sweeney, D. Gillett, D. Evseev, D. Calabrese, H. Hoppe, A. Kirk, and S. Sullivan, "High-quality streamable free-viewpoint video," *ACM Trans. Graph.*, vol. 34, no. 4, 2015.
- [18] C. Chi and S. Song, "Garmentnets: Category-level pose estimation for garments via canonical space shape completion," in *The IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [19] Z. Huang, X. Lin, and D. Held, "Mesh-based Dynamics with Occlusion Reasoning for Cloth Manipulation," in *Proceedings of Robotics: Science and Systems*, New York City, NY, USA, June 2022.
- [20] W. Wang, G. Li, M. Zamora, and S. Coros, "Trtm: Template-based reconstruction and target-oriented manipulation of crumpled cloths," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 12 522–12 528.
- [21] S. Qian, K. Mo, V. Blukis, D. F. Fouhey, D. Fox, and A. Goyal, "3d-mvp: 3d multiview pretraining for robotic manipulation," 2025. [Online]. Available: <https://arxiv.org/abs/2406.18158>
- [22] W. Agnew, C. Xie, A. Walsman, O. Murad, Y. Wang, P. Domingos, and S. Srinivasa, "Amodal 3d reconstruction for robotic manipulation via stability and connectivity," in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, vol. 155. PMLR, 2021, pp. 1498–1508.
- [23] E. Amin Mansour, H. Zheng, and R. K. Katzschmann, "Fast point cloud to mesh reconstruction for deformable object tracking," in *International Conference on Robotics, Computer Vision and Intelligent Systems*. Springer, 2024, pp. 391–409.
- [24] J. Lei and K. Daniilidis, "Cadex: Learning canonical deformation coordinate space for dynamic surface representation via neural homeomorphism," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 6624–6634.
- [25] M. Niemeyer, L. Mescheder, M. Oechsle, and A. Geiger, "Occupancy flow: 4d reconstruction by learning particle dynamics," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [26] N. Wang, Y. Zhang, Z. Li, Y. Fu, H. Yu, W. Liu, X. Xue, and Y.-G. Jiang, "Pixel2mesh: 3d mesh model generation via image guided deformation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 10, pp. 3600–3613, 2021.
- [27] D. Jack, J. K. Pontes, S. Sridharan, C. Fookes, S. Shirazi, F. Maire, and A. Eriksson, "Learning free-form deformations for 3d object reconstruction," in *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part II 14*. Springer, 2019, pp. 317–333.
- [28] A. Kanazawa, S. Tulsiani, A. A. Efros, and J. Malik, "Learning category-specific mesh reconstruction from image collections," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 371–386.
- [29] J. Xu, W. Cheng, Y. Gao, X. Wang, S. Gao, and Y. Shan, "Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models," *arXiv preprint arXiv:2404.07191*, 2024.
- [30] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby *et al.*, "DinoV2: Learning robust visual features without supervision," *arXiv preprint arXiv:2304.07193*, 2023.
- [31] Y. Yang, C. Feng, Y. Shen, and D. Tian, "Foldingnet: Point cloud auto-encoder via deep grid deformation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 206–215.
- [32] G. Turk, "The stanford bunny," The Stanford Graphics Laboratory, Tech. Rep., 1994.
- [33] N. Noor, A. Shapira, R. Edri, I. Gal, L. Wertheim, and T. Dvir, "3d printing of personalized thick and perfusable cardiac patches and hearts," *Advanced science*, vol. 6, no. 11, p. 1900344, 2019.
- [34] S. Peng, C. Jiang, Y. Liao, M. Niemeyer, M. Pollefeys, and A. Geiger, "Shape as points: A differentiable poisson solver," *Advances in Neural Information Processing Systems*, vol. 34, pp. 13 032–13 044, 2021.
- [35] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '96)*, 1996, pp. 303–312.
- [36] Y. Zheng, Q. Zhao, G. Yang, W. Yifan, D. Xiang, F. Dubost, D. Lagun, T. Beeler, F. Tombari, L. Guibas, and G. Wetzstein, "Physavatar: Learning the physics of dressed 3d avatars from visual observations," in *Computer Vision – ECCV 2024*, 2025, pp. 262–284.
- [37] H. Chen, D. Zhao, and J. Barbič, "Capturing animation-ready isotropic materials using systematic poking," *ACM Trans. Graph.*, vol. 42, no. 6, Dec. 2023.
- [38] NVIDIA, "Isaac sim," NVIDIA Developer Forum, 2021. [Online]. Available: <https://developer.nvidia.com/isaac/sim>
- [39] M. Geilinger, D. Hahn, J. Zehnder, M. Bächer, B. Thomaszewski, and S. Coros, "Add: analytically differentiable dynamics for multi-body systems with frictional contact," *ACM Trans. Graph.*, vol. 39, no. 6, Nov. 2020.
- [40] E. Heiden, M. Macklin, Y. S. Narang, D. Fox, A. Garg, and F. Ramos, "DiSECT: A Differentiable Simulation Engine for Autonomous Robotic Cutting," in *Proceedings of Robotics: Science and Systems*, Virtual, July 2021.
- [41] P. Ma, P. Y. Chen, B. Deng, J. B. Tenenbaum, T. Du, C. Gan, and W. Matusik, "Learning neural constitutive laws from motion observations for generalizable pde dynamics," in *International Conference on Machine Learning*. PMLR, 2023, pp. 23 279–23 300.
- [42] K. R. Allen, T. L. Guevara, Y. Rubanova, K. Stachenfeld, A. Sanchez-Gonzalez, P. Battaglia, and T. Pfaff, "Graph network simulators can learn discontinuous, rigid contact dynamics," in *Conference on Robot Learning*. PMLR, 2023, pp. 1157–1167.
- [43] Formlabs, "Guide to transparent 3d printing," Formlabs blog, 2025. [Online]. Available: <https://formlabs.com/blog/3d-printing-transparent-parts-techniques-for-finishing-clear-resin/>
- [44] C. C. Christoph, M. Eberlein, F. Katsimalis, A. Roberti, A. Sympetheros, M. R. Vogt, D. Liconti, C. Yang, B. G. Cangan, R. J. Hinchet *et al.*, "Orca: An open-source, reliable, cost-effective, anthropomorphic robotic hand for uninterrupted dexterous task learning," *arXiv preprint arXiv:2504.04259*, 2025.