

# More with LESS – Local Scene Representations for Tactile Imaging

Zohar Rimon<sup>\*,1</sup> Elisei Shafer<sup>1</sup> Tal Tepper<sup>1</sup> Daniel Kozin<sup>1</sup> Alon Malka<sup>1</sup>  
 Roy Holland<sup>1,2</sup> Aviv Tamar<sup>1</sup>

<sup>1</sup>Technion - Israel Institute of Technology <sup>2</sup>Rambam Health Care Campus

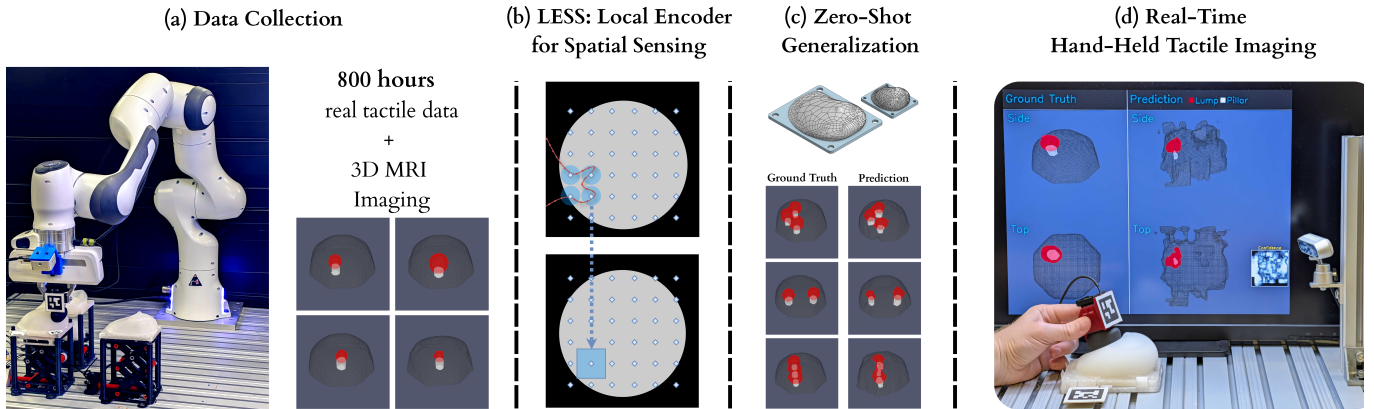


Fig. 1: **Overview of our approach:** (a) Using an automated robotic data collection setup, we collect the largest soft-body tactile interaction dataset to date, coupled with MRI imaging of the inner structure of our modular-designed phantoms. (b) We propose the LESS neural network architecture to effectively represent a complex tactile scene by capturing local tactile interaction sequences, and generate local visualizations of the sensed objects. (c) Our approach leads to zero-shot compositional generalization and maintains accurate predictions on unseen phantom configurations with multiple inclusions and different shapes. (d) Finally, we show a proof-of-concept real-time hand-held tactile imaging system that visualizes the inner structure of soft objects.

**Abstract**— Tactile imaging seeks to reconstruct the internal structure of soft objects through touch sensing, with applications in medical diagnosis and robotic manipulation. Recent self-supervised learning approaches have shown promising results, but rely on global, unstructured representations and robot-controlled sensing, limiting generalization and practical use. We propose Local Encoder for Spatial Sensing (LESS), an object-centric tactile representation that exploits the local nature of touch. The tactile scene is modeled as a grid of recurrent encoders with local receptive fields, whose states are fused to reconstruct 2D or 3D images of internal structure. This compositional design enables strong generalization: models trained on single-inclusion phantoms accurately image objects with multiple inclusions and varying sizes. The local structure further supports spatial uncertainty estimation. In addition, we enable hand-held tactile imaging via external pose tracking and human-like palpation data, and extend tactile imaging to full 3D reconstruction.

## I. INTRODUCTION

Humans have a remarkable capability to understand and visualize the world through touch. Tactile scene understanding is important for medical tactile imaging (TI, [9]) and for unlocking robotic capabilities in medical, domestic, and agricul-

tural domains. Although tactile sensor hardware has advanced recently [21, 22, 43, 8], scene understanding from tactile data is still a challenging algorithmic problem, especially when non-rigid objects are involved.

The recent work of Rimon et al. [34] proposed a self-supervised learning (SSL) based approach to TI. In their work, a manipulator equipped with a tactile sensor automatically palpates a set of soft objects (breast phantoms), for which ground truth internal structure is available through magnetic resonance imaging (MRI). A tactile representation is learned from the SSL data by training an encoder-decoder model. The encoder processes a sequence of tactile measurements and sensor positions into a “representation” vector, while the decoder predicts from the representation the tactile measurement at a future position. Following the SSL training, an image generation model is trained to predict internal structure (MRI slice) from the tactile representation. Rimon et al. [34] showed that with sufficient SSL training, the network can learn complex patterns in the tactile data to yield accurate imaging.

However, several features in the method of Rimon et al. [34] limit its practical applicability for tactile imaging. The first is the representation. Representing a tactile scene using a single *unstructured* vector ignores the compositional nature of physical objects, and requires training on “all possible object

\* Correspondence to zohar.rimon@campus.technion.ac.il.  
 Website: [zoharri.github.io/LESS](https://zoharri.github.io/LESS)  
 Code: [github.com/zoharri/LESS/](https://github.com/zoharri/LESS/)  
 Data: [zenodo.org/communities/artificial-palpation](https://zenodo.org/communities/artificial-palpation)

combinations” to generalize, which may be too costly. Our insight, similar to object-centric representations in computer vision [27, 7], is that touch is inherently local: the tactile response of the internal structure of an object at some position is *nearly independent* of the internal structure at far away positions. Based on this idea, we propose **Local Encoder for Spatial Sensing (LESS)** – a novel object-centric representation for tactile scenes. The LESS encoder is composed of a grid of ‘particles’<sup>1</sup>, where each particle is a recurrent neural network that processes the sequence of tactile measurements within its spatial receptive field. The LESS image generator combines the latent state from all the particles into a single 2D or 3D image of the scene. We show that by training on phantoms with single inclusions as in [34], we can produce accurate tactile images for objects with multiple inclusions at test time. Moreover, we demonstrate that LESS yields accurate imaging for objects of different size from training. Finally, we show that the local structure of LESS allows us to visualize the *uncertainty* of the tactile scene in different areas, which may be used to direct the operator to collect additional measurements.

Another shortcoming in the method of Rimon et al. [34] is that it must be robot-operated at test time, limiting its use as a hand-held tactile imaging device. The reason is twofold: the method requires accurate pose tracking of the sensor obtained by the robot kinematics, and the training data contains only vertical movements, which are very different from natural human palpation. We relax these limitations by adding external position tracking of the sensor, and collecting robot palpation data that more resembles human motion. These technical modifications allow us to demonstrate the first hand-held tactile imaging device based on SSL. Finally, the imaging in [34] was limited to 2D. We extend the method to 3D imaging, and interestingly find that 3D reconstruction also improves upon the results in [34] when considering 2D image slices from our predicted 3D volume.

## II. RELATED WORKS

**Tactile imaging:** Mechanical/tactile imaging systems formalized pressure–position sensing and inversion to characterize tissue stiffness and lesion geometry in vivo [45, 9, 10, 44, 11]. Comprehensive reviews position tactile imaging as a low-cost elastography modality [37] and outline requirements for broader diagnostic adoption [38, 16]. However, state-of-the-art methods are based on averaging readings of tactile sensor arrays to generate a force map, and cannot detect subtle tissue features that are apparent to human touch and important for diagnosis [9, 35]. In breast imaging, Hampson et al. [16] highlight the inability to detect background tissue elasticity, lesion position within the breast, and lesion acutance and mobility – capabilities that an SSL approach can potentially acquire with enough training data. A related line of work explored robotic palpation [19, 18, 42, 39, 36], while we only

<sup>1</sup>In this work, we use the term particle to describe a neural network representation that is also situated in space, inspired by the object-centric representation literature [6]. This term should not be confused with particle filters from the state estimation literature.

use a robot for data collection; as we demonstrate, our artificial palpation device can be hand-held.

**Representations for tactile data:** Neural networks were used for various tactile processing tasks, including material classification [14, 47, 28, 1], grasp-success prediction [3], slip detection [24], and contour-following control [23]. Representation learning for tactile data is an emerging area.

Several studies investigated learning representations for image-based tactile sensors. Sparsh [17] is an SSL representation for single tactile measurements, and AnyTouch [13] is an SSL representation for short (3-frame) tactile videos. For tactile sensor arrays, T-DEX [15] is an SSL representation for a single measurement of a taxel grid. All of these methods learn to represent individual tactile measurements, while our work is on representing tactile *scenes*.

For tactile scene representations, prior work excluding Rimon et al. [34] focused almost exclusively on rigid objects. Qi et al. [32] learned a representation for a sequence of tactile, proprioceptive, and visual observations, using simulated 3D point-clouds of objects during training. NeuralFeels [41] uses neural fields to represent the 3D shape of an object from vision and tactile sensors, Tactile-Informed 3DGS [5] employ Gaussian splatting for similarly structured data, and VTacO [46] predict the surface of an object from vision and touch by learning a winding number field. None of these methods can be used for imaging the inner structure of soft objects, which is critical for medical applications and our focus here.

## III. FORMULATION

We follow the artificial palpation problem formulation of Rimon et al. [34]. A rigid tactile sensor is controlled to be at time  $t \in 0, \dots, T$ , in pose  $x_t \in \mathbb{R}^6$ . The sensor touches a soft body  $M$ , and produces a  $k$ -dimensional force reading  $f_t \in \mathbb{R}^k$ . Here,  $M$  is generally unknown and represents all the structural and mechanical properties of the body that determine the force on the sensor. Subsequently, the sensor is moved to the next pose  $x_{t+1}$  by a controller, and the next reading is obtained. In addition, we may have an observation of the soft body, denoted  $I$ , for example, an MRI scan.

We make two modifications to the formulation in [34]. The first is sensor pose estimation: we cannot observe  $x_t$ , but only a noisy estimate  $\hat{x}_t$ . The second is the controller. While [34] considered a fixed motion for the sensor, we investigate both different controllers and human-operated hand-held control. We denote  $x_{t+1} = g(h_t)$  a general control function that depends on  $h_t$ , the history of measurements up to time  $t$ .

## IV. METHOD

This section describes our method for learning tactile imaging. We begin with the LESS architecture, continue with training details and data collection, and finally describe technical contributions required for hand-held tactile imaging.

### A. LESS - Local Encoder for Spatial Sensing

The method of [34] proposed a recurrent neural network that encodes a sequence of tactile measurements

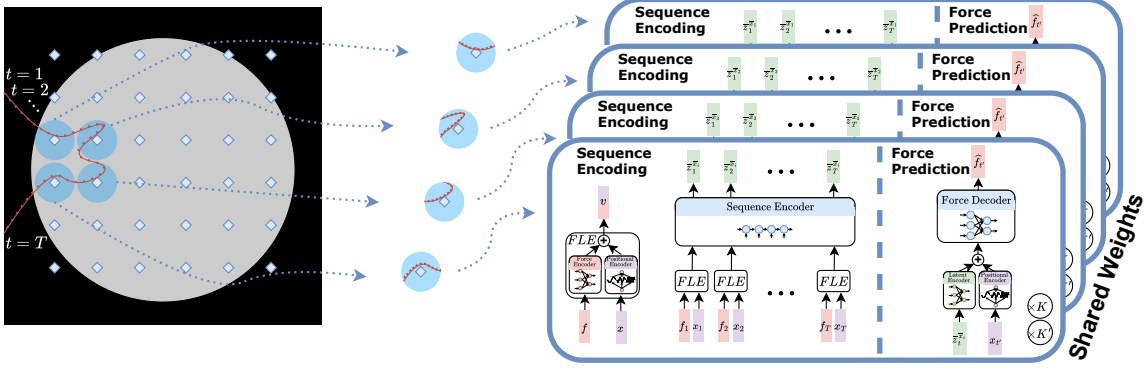


Fig. 2: **LESS representation learning:** A set of localized representations for processing a tactile scene, where each representation attends only to measurements within its receptive field. For a localized representation positioned at  $\bar{x}$ , the sequence of tactile data within its receptive field is centered around  $\bar{x}$ , and processed with a GRU-based encoder-decoder architecture, with shared weights. The encoder-decoder figure on the right is modified from [34].

$\{x_0, f_0, \dots, x_t, f_t\}$  into a  $d_z$ -dimensional latent vector representation  $z_t \in \mathbb{R}^{d_z}$ . Our main insight is that tactile scenes, where material properties at a particular position  $x$  typically do not affect the force sensed at a distant position  $x'$ , can be better represented by a collection of  $N$  local representations  $\bar{z}_t = \{z_t^{\bar{x}_1}, \dots, z_t^{\bar{x}_N}\}$ , where each local representation  $z_t^{\bar{x}_i}$  has a *spatial receptive field* centred around position  $\bar{x}_i$ . The motivation for such a disentangled representation is two-fold, as we describe next.

**Compositional Generalization:** Consider objects that are composed of several different material properties at different positions. A single vector representation must learn to represent *all the combinations* of different materials at different positions, the quantity of which grows exponentially with the number of materials. Our disentangled representation only needs to learn the local tactile response for the possible materials, while the compositional generalization is by construction.

**Localized Uncertainty Estimates:** For a hand-held device, it is important to convey to the operator areas with high uncertainty, to guide data collection at test time. We hypothesize that localizing the uncertainty using a disentangled representation is easier than using a single vector representation, where uncertainty at multiple positions may be entangled.

Our architecture, **Local Encoder for Spatial Sensing (LESS)**, is composed of a tactile sequence encoder-decoder module, and an image reconstruction module, depicted in Figure 2 and Figure 3, and described in the following.

1) **LESS Sequence Encoder:** The local representation  $z_t^{\bar{x}_i}$  is obtained from the input sequence  $\{x_0, f_0, \dots, x_t, f_t\}$  by applying the following sequence of steps.

**Receptive Field:** Each  $\bar{x}_i$  has a receptive field  $\mathcal{X}_i = \{x : D(x, \bar{x}_i) < \beta\}$ , for some distance function  $D$ . We filter out inputs that are outside the receptive field of  $\bar{x}_i$ , that is, only pairs  $x_j, f_j$  that satisfy  $x_j \in \mathcal{X}_i$  are kept in the input sequence to  $\bar{x}_i$ .

**Centering:** Every input pose  $x_j$  in the input sequence to  $\bar{x}_i$  is centered around  $\bar{x}_i$  by subtracting  $\bar{x}_i$  from it.

**Force-Location Encoder (FLE):** We follow [34] and encode each input pair  $x_j, f_j$  by first encoding the centered  $x_j$  using sinusoidal positional encoding (PE), then applying an MLP on  $f_j$ , and finally summing the MLP and PE outputs.

**Recurrent Neural Network:** The sequence of FLE encoded forces and poses is input to a Gated Recurrent Unit (GRU, [4]), and  $z_t^{\bar{x}_i}$  is obtained as the GRU hidden state.

In our implementation, the  $N$  GRUs for different positions have shared weights, as relevant patterns in localized sensory readings from different positions should be similar.

In addition, we positioned the  $\bar{x}_i$ 's on a uniform 2-dimensional grid, and our distance function was Euclidean on the 2-dimensional  $x - y$  plane, namely,  $D(a, b) = \sqrt{(a_x - b_x)^2 + (a_y - b_y)^2}$ . This choice is motivated by the fact that our objects mostly vary in the planar position of different inclusions. More general settings may require a 3-dimensional grid, and possibly accounting for orientations.

**Force Decoder:** To predict  $\hat{f}(x, z_t^{\bar{x}_i})$ , the force at position  $x$  from a local representation  $z_t^{\bar{x}_i}$ , we first center  $x$  around  $\bar{x}_i$ , as described above, then apply PE to the centered  $x$  and an MLP to  $z_t^{\bar{x}_i}$ , and finally input the sum of the PE and MLP to another MLP that predicts the force.

2) **LESS Image Reconstruction:** Given the encoded representation, which essentially captures information about the palpated object, we seek to produce an image of the object's internal structure. In [34], this was done by mapping the single representation vector  $z_t$  to a 2-dimensional image using transposed convolutions.

In contrast, we seek to map the set of local representations  $\{z_t^{\bar{x}_1}, \dots, z_t^{\bar{x}_N}\}$  to an image such that each reconstructed pixel is only affected by representations in its local area. In addition, we seek to image 3-dimensional volumes, to capture a richer and more precise visualization of objects' internal structure. We begin by describing our approach to 2-dimensional images, and later extend to 3D.

Consider a square image patch centered on the  $x - y$  components of  $\bar{x}$ , with side length  $2\rho$ . We assume that each




Fig. 3: **LESS image reconstruction:** Each learned local representation is mapped to a local image patch using a transposed convolution neural network. In turn, all patches are stitched together to produce the final image by averaging logits and applying a softmax.

patch contains  $N_p \times N_p$  pixels, and that each pixel can take  $\kappa$  possible values; in our experiments,  $\kappa = 3$ , corresponding to free space, hydrogel, or silicone inclusion. Let  $p(\bar{x}) \in \kappa^{N_p \times N_p}$  denote the logits for every pixel in  $\bar{x}$ 's image patch. The logits for each patch are obtained using a transposed convolution with input  $\bar{z}^{\bar{x}}$ , where the transposed convolutions for all patches have shared weights. To compose the patches into a complete image, we first sum all logits and then apply a softmax to obtain class probabilities for every pixel.

**3-dimensional reconstruction** The 3-dimensional reconstruction module is similar, with the difference that each patch is a square column with  $N_p \times N_p \times N_z$  pixels.

**Related Spatio-Temporal Architectures:** The LESS architecture draws inspiration from related spatio-temporal representations in the literature, such as the convolutional LSTM [40] and Deep Latent Particles [6]. Several properties specialized for tactile imaging differentiate it from previous work. First, different from convolutional neural networks, LESS particles apply a hard threshold outside their receptive fields, which relates to the physical locality of touch. Second, LESS decodes both forces and 2D/3D images.

## B. SSL Training and Data

Following [34], force prediction drives SSL of the LESS encoder-decoder. However, we extend the force prediction method of [34] to a localized setting.

Our training loss  $\ell = \frac{1}{N} \sum_{i=1}^N \bar{\ell}(\bar{x}_i)$  averages the loss functions for individual particles, which we describe next. Consider a local representation centered on  $\bar{x}_i$ , and a sequence  $\{x_0, f_0, \dots, x_T, f_T\}$  in the data. We first filter out pose-force pairs outside  $\mathcal{X}_i$ , and calculate the local representations  $\{\bar{z}_t^{\bar{x}_i}\}$  as described above for all time steps in the filtered sequence. The mean-squared error (MSE) in predicting a force at position  $x_{t'}$  given the representation  $\{\bar{z}_t^{\bar{x}_i}\}$  is  $\|\hat{f}(x_{t'}, \bar{z}_t^{\bar{x}_i}) - f_{t'}\|^2$ . The loss  $\bar{\ell}(\bar{x}_i)$  averages the MSE over a set of time index pairs  $t, t'$  sampled similarly to [34] (full details in Appendix E).

SSL data consists of force and pose sequences obtained from palpating soft phantoms in the training set. We collect a similar dataset to the Artificial Palpation Poke Dataset [33], which was also used in [34], and henceforth termed `data-v1`. The sequences in this data correspond to vertical

motions of the tactile sensor in a fixed orientation, touching artificial breast phantoms made from soft silicone skin and hydrogel filling, with a single, round, soft-silicone inclusion. We hypothesize that hand-held motion would be significantly different, and therefore produced two new training datasets, `data-poke` and `data-primitive`, and a small test dataset `data-handheld`. Our data generation method differs from [34] in object fabrication method, robot palpation motion, and the automatic data collection mechanism, as we describe next; full details and reproduction steps are in Appendix A. All datasets are publicly available.

**Phantom fabrication:** Our fabrication technique is based on curing soft silicone in 3D printed moulds, which enables accurate thickness control over the phantom 'skin', and inclusions with complex shapes, such as several connected ball-shaped objects. Following the high-level design in [34] of an outer shell and an insert with an inclusion that can be rotated in 8 angles, we generated 40 inserts (24 of which have been proposed in [34]) and 8 shells, with inclusions varying in size and shape. See Appendix A for details.

**Tactile Sensor:** We use an off-the-shelf Xela uSCu ALHA sensor [43], capturing data at 100Hz. We emphasize that our learning method is not sensor-specific.

**Sensor Motion:** `data-poke` was collected using poke trajectories, similarly to [34], but with various orientations of the sensor, which are fixed throughout the motion. For data that better resembles human motion, we recorded several motion primitives using teleoperation, and tracked them using position control to generate `data-primitive`. Finally, test `data-handheld` was generated by manual motion, with poses recorded using a fiducial-based pose tracking device (see Appendices F and I for details). Following [34], we subsample the sensor measurements 1 : 50 in `data-poke`. For `data-primitive`, we subsampled the measurements 1 : 10, which we found to work better due to more complex motion. Each phantom in `data-poke` and `data-primitive` was poked using 100 trajectories, while in `data-handheld`, with 70. For all datasets, we concatenate the trajectories as input to LESS.

**Automatic Data Collection:** To speed up data collection, we designed a motorized scissor lift to automatically reorient an insert within its shell as seen in Figures 4a and 4b (full details are in Appendix H). Using 3 lifts and a Franka Panda manipulator, we collected data almost continuously for  $\sim 30$  days, resulting in  $\sim 800$  hours of tactile data, over  $7 \times$  the size of `data-v1`.

## C. Tactile Imaging Training and Data

The output of the LESS image reconstruction model is either a 2D or 3D image. Similarly to [34], we train the model to predict images of the internal structure of the insert being palpated, obtained using MRI, as available in `data-v1` [33]. A straightforward loss function for image reconstruction is cross entropy between the predicted image and the MRI ground truth. However, especially for 3D images, we found that our data suffers from severe class imbalance, as the

inclusion is relatively small compared to the gel filling of the insert. We therefore opted for a combination of the Focal loss [25] and Dice loss [29], which we found to empirically improve performance, full details are in Appendix D.

Ground truth MRI for the inserts we produced was obtained using a 3T MRI (Siemens Prisma) system with a 64-channel coil, using the protocol described in [34]. Different from [34], which only considered a 2-dimensional horizontal slice of the scan, we consider the full 3D volume. Full details on preparing the ground truth images are in Appendix B.

#### D. Hand-held Tactile Imaging

Considering applications of tactile imaging, the hand-held form factor is essential for clinical usage, as it allows the physician to maneuver and maintain patient contact during examination, similar to the workflow in ultrasound imaging. It may also be important for home or bedside devices. Next, we detail a proof-of-concept hand-held tactile imaging system.

**Pose Estimation and Generalization** Our approach requires the sensor pose at each tactile measurement. In practice, sensor pose must be estimated, and we evaluated the robustness of our imaging to sensor pose noise. We found that when noise is i.i.d. Gaussian, our method remains stable to errors up to  $\sim 1$  mm in position and  $\sim 0.1$  radian in orientation. Based on these results, we designed a fiducial-based pose estimation device with average pose error that is smaller than the stability thresholds above, as detailed in Appendix F. Using our system, we can palpate general objects manually to collect test data. Further, to generalize imaging to objects with different dimensions from the training objects, we extend the set of particles in the image decoder to include particles that cover a larger volume at test time. Since particles share weights, we can add an arbitrary number of particles, so long as we maintain a similar spatial resolution to the training data.

**Uncertainty Estimation Guidance** Hand-held imaging devices such as ultrasound probes are known to be *operator dependent*, a generally undesirable quality [31]. To mitigate operator dependence, we propose to provide online feedback on imaging uncertainty, potentially guiding the operator to collect additional samples in uncertain regions. The LESS architecture naturally accommodates *local* uncertainty estimates: for each pixel in the image, we measure uncertainty using the Shannon entropy of the class probabilities. Since each pixel is affected only by local representations, the uncertainty is only affected by samples collected in each pixel’s vicinity.

## V. EXPERIMENTS

We conduct a series of experiments to study the following. (1) Qualitative and quantitative assessment of 3D tactile imaging using our learning based approach and datasets. (2) Can LESS yield zero-shot compositional generalization? (3) Is LESS effective for local uncertainty estimation? and (4) Can our system and data be used to demonstrate a proof-of-concept real-time hand-held tactile imaging?




Fig. 4: (a) and (b) Insert being rotated and pushed into the shell in the different states of the scissor lift (c) 3D ground truth for two inserts after post-processing of its MRI scan. Blue: background. Gray: insert. Red: inclusion. White: supporting pillar.

TABLE I: Comparison of global representation training on 2D and 3D labels. The first three rows are for 2D slice metrics and the rest are 3D metrics.

	2D Training	3D Training	Average Pred.	
2D	$\Delta_{\text{Diameter}}$ [mm] ↓	$2.7 \pm 0.1$	$1.8 \pm 0.0$	2.02
	$\Delta_{\text{CoM}}$ [mm] ↓	$2.1 \pm 0.1$	$1.7 \pm 0.1$	12.1
	F1 Score [%] ↑	$79.9 \pm 0.8$	$85.2 \pm 0.8$	–
3D	$\Delta_{\text{Diameter}}$ [mm] ↓	–	$0.8 \pm 0.0$	1.44
	$\Delta_{\text{CoM}}$ [mm] ↓	–	$1.7 \pm 0.1$	12.09
	F1 Score [%] ↑	–	$83.5 \pm 0.7$	–

Throughout all experiments, we report the standard deviation by running 3 random seeds over train/test splits and random weights initialization.

#### A. 3D Tactile Imaging

As discussed in IV-A, we extend [34] to 3D imaging allowing for richer visualization of the size, shape, and position of internal inclusions.

An example of our 3D ground-truth is shown in Figure 4c. Our inserts (which are positioned inside breast-shaped shells) are filled with gel, with an outer skin made of soft silicone. In the insert, we place a round inclusion connected by a narrow pillar to the base of the insert, both made out of soft silicone. For visual clarity, we colour the pillar differently from the round inclusion, using an ad-hoc image processing method, described in Appendix B. The pillar is not touched by the sensor during poke motions; yet, it can be reconstructed, since the method was trained on ground truth images that include the pillars. This is a demonstration of the learned *inference* capabilities of our approach. In the following, we quantitatively compare the accuracy of our new imaging results with the 2D imaging technique of [34].

We begin by comparing our 3D reconstruction approach with the state-of-the-art approach of [34], and study how the dimensionality of the data affects reconstruction quality.




Fig. 5: Prediction of LESS for intermediate steps of four input sequences (steps were chosen to demonstrate development of the image over time). The ground-truth is provided on the right side. Different rows show: (1) single inclusion, (2) two inclusions, (3) three connected inclusions, (4) larger inclusion. As described in the text, the model was trained on single inclusions only.

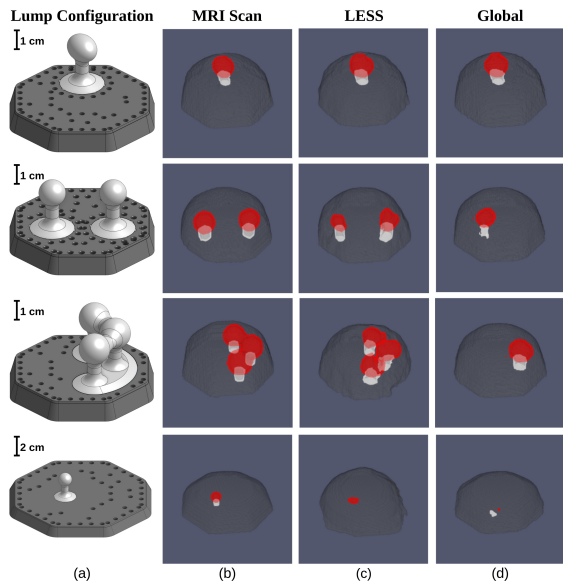


Fig. 6: Comparison between multiple samples from our in-distribution (first row) and out-of-distribution (last three rows) inserts. See text for details. Columns show: (a) 3D CAD (b) 3D MRI ground-truth (c) image reconstruction using LESS (d) image reconstruction using GLOBAL.

To separate the contribution of the data dimensionality from other architectural improvements in LESS, we trained the representation of [34] on `data-poke` with 2D and 3D MRI labels, where the 3D training is based on our method in Section IV-C.

For the 2D labels, we report the F1 score, inclusion Center-of-Mass (CoM) error, and inclusion diameter error. For the 3D labels we report these metrics on a constant height slice and the 3D version of these metrics. Exact definition of the

metrics can be found in Appendix C. To connect our results to easily interpretable quantities, we also report the metrics for predicting the average diameter and CoM over the dataset.

Surprisingly, Table I shows that incorporating 3D predictions leads to consistent improvements in 2D metrics compared to the results reported in [34]. We hypothesize that this improvement originates from the additional geometric prior provided by the 3D representation. When sufficient training data is available, learning from volumetric information encourages the model to capture spatial regularities that are not observable in isolated 2D slices, such as consistency across adjacent slices. Thus, in addition to producing more interpretable visualizations, 3D data also yields more accurate reconstructions.

For the rest of the experiments, we train all models with the 3D labels, but report the 2D metrics, which are easier to interpret, and we term the global baseline of [34], trained on the 3D labels of `data-poke` as GLOBAL.

### B. Zero-Shot Compositional Generalization with LESS

One of the core benefits of LESS is the emergent compositional generalization (as discussed in Section IV-A). To demonstrate and evaluate this, we manufactured a set of specialized phantoms, unseen during training, that test compositional generalization to changes in the shape of the object and the number of inclusions inside it. Accordingly, we split the specialized phantoms into two categories - `multiple` and `large`. The `multiple` set corresponds to phantoms with multiple inclusions in a single insert, and the `large` set corresponds to a phantom with the same height, but a base with a  $4\times$  larger area than the regular phantoms, and with single inclusions with similar shapes and sizes to the regular phantoms. Full details of the specialized phantoms can be found in Appendix A. We emphasize that in this section, we trained models on the `data-poke` dataset, which only




Fig. 7: Prediction and confidence of LESS compared to GLOBAL for intermediate steps of the same input sequence.

includes phantoms with *single* inclusions. Thus, phantoms with *multiple* are essentially out of distribution. We collected test interactions on the specialized phantoms with the same controller as in `data-poke`.

TABLE II: Results of GLOBAL and LESS on in-distribution single inclusion and out-of-distribution multiple inclusions and large phantom. In the *multiple* set, we report the total area error instead of the diameter error. Due to their different inclusion composition, the metrics of different insert sets are not comparable. We clarify that in the *large* set, the diameter and CoM errors are expected to be larger.

		GLOBAL	LESS (ours)
Single	$\Delta_{\text{Diameter}}$ [mm] ↓	$1.8 \pm 0.0$	$1.5 \pm 0.1$
	$\Delta_{\text{CoM}}$ [mm] ↓	$1.7 \pm 0.1$	$1.7 \pm 0.1$
	F1 Score [%] ↑	$85.2 \pm 0.8$	$83.5 \pm 0.6$
Multiple	$\Delta_{\text{Area}}$ [mm <sup>2</sup> ] ↓	$78.4 \pm 5.1$	$46.1 \pm 2.9$
	$\Delta_{\text{CoM}}$ [mm] ↓	$7.4 \pm 0.4$	$4.0 \pm 0.0$
	F1 Score [%] ↑	$44.6 \pm 1.3$	$71.1 \pm 0.0$
Large	$\Delta_{\text{Diameter}}$ [mm] ↓	$9.0 \pm 1.8$	$2.7 \pm 1.1$
	$\Delta_{\text{CoM}}$ [mm] ↓	$26.8 \pm 2.8$	$2.3 \pm 0.2$
	F1 Score [%] ↑	$1.1 \pm 0.5$	$71.3 \pm 3.6$

We trained LESS on `data-poke` similarly to GLOBAL. LESS is able to effectively reconstruct complex internal structure as highlighted in Figure 5. We compare LESS and GLOBAL on other types of phantoms at test time. As seen in Table II, LESS performs slightly worse compared to GLOBAL on in-distribution single inclusions, since each image patch that LESS generates is affected by less information (due to the receptive fields), making the prediction problem more difficult. However, LESS significantly outperforms the baseline on the out-of-distribution benchmark.

As shown in Figure 6, GLOBAL does not generalize to

the *multiple* set, and generates an image with only one of the inclusions. This is expected, as training on single inclusion data is not sufficient for generalization to more complex inner structure when using global representations. On the other hand, LESS is able to generalize the local patterns of single inclusions in the data to accurately generate multiple inclusions. The *large* set is completely out of distribution for GLOBAL (as it was never trained on poses that are encountered in this set). On the other hand, due to LESS’s centering process, it is able to generalize to this set as well. We believe these results show the potential of LESS for tactile imaging of generally structured objects.




Fig. 8: LESS performance for different receptive field sizes.

### C. Receptive Field Tuning

An important hyperparameter of LESS is the receptive field size  $\beta$ . As can be seen in Figure 8, increasing the receptive field allows more information to be encoded in each particle, improving the performance on the *single* set and eventually achieving similar accuracy to GLOBAL. The tradeoff in increasing  $\beta$  is reducing the local invariance, thereby reducing compositional generalization, as evident by the degraded performance on the *multiple* set in Figure 8. Finally, we observe that  $\beta \approx 10[\text{mm}]$  results in the best

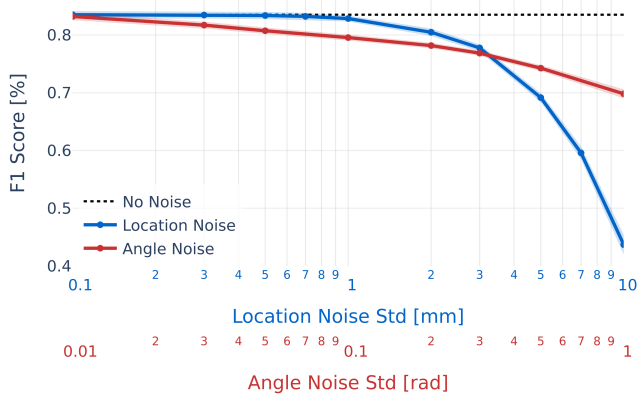


Fig. 9: Effect of pose noise on the F1 score of LESS.

performance on the multiple set, without compromising accuracy on the single set; we used this  $\beta$  value throughout our experiments.

#### D. Local Uncertainty Estimation

We qualitatively evaluate uncertainty estimation using LESS and compare with GLOBAL. For both methods we use the Shannon entropy for pixel level uncertainty, as described in Section IV-D; the only difference is in the neural network that generates pixel probabilities. In Figure 7 we show the image reconstruction and uncertainty of both models during intermediate steps of reconstructing an image from the same tactile input sequence. As can be seen, LESS produces a significantly more informative uncertainty estimation, showing low confidence on unobserved areas, while GLOBAL collapses to a very high confidence early on, before many areas of the phantom have been touched. To further understand these results, note that GLOBAL yields high confidence in areas where it lacks sufficient data, such as those outside the insert area. The reason is that these pixels are always outside the insert area in the training data. In comparison, LESS produces very low confidence in these coordinates, until it receives sufficient data in their vicinity.

#### E. Hand-Held Tactile Imaging

Finally, we investigate whether our tactile imaging system can be used as a hand-held device that produces accurate imaging in real-time. The two main challenges in this task are tracking the pose of the tactile sensor accurately, and the distribution shift between sensor motion in the robot-operated data collection and manual hand-held motion. In this section we first seek to test the robustness of LESS to inaccurate pose, which is crucial for certain design choices in our pose estimation system. Then, we show that by training on data-primitive we substantially reduce the distribution shift that occurs when testing on human motions.

**Robustness to Inaccurate Pose** To determine the required system accuracy, we evaluated the sensitivity of the model to errors in the sensor pose. Two types of noise were injected: translational and angular.  $x_t \in \mathbb{R}^6$  is the true pose, where the

first 3 components are  $x$ - $y$ - $z$  positions, and the last 3 components are pitch-roll-yaw angles. We modeled the noise as additive i.i.d. zero-mean Gaussian noise:  $\hat{x}_t = x_t + \mathcal{N}(0, \text{diag}(\bar{\sigma}))$ , where  $\bar{\sigma} = [\sigma_{pos}, \sigma_{pos}, \sigma_{pos}, \sigma_{angle}, \sigma_{angle}, \sigma_{angle}]$ . We evaluated tactile imaging results with varying noise in the test input poses, shown in Figure 9. Based on these results, we have designed our pose estimation system as discussed in Section IV-D. We verify the noise level of our pose to be less than  $1\text{mm}$  and  $0.02\text{rad}$  and further measure end-to-end performance degradation due to pose error in Appendix G.

**Training on More Diverse Data** To reduce the distribution shift when testing on human motions in the hand-held setup, we propose to train LESS on data-primitive. This dataset contains a much more diverse set of motions compared to the constant angle poke trajectories in data-poke. We found that combining data-poke with data-primitive as the training data yielded the best results. An ablation study on the training data is presented in Table III.

TABLE III: Training LESS on data-primitive. We trained on only data-poke (**Poke**), only data-primitive (**Pri**) and the union of data-poke, and data-primitive (**Pri + Poke**). All methods were tested on data-primitive only.

Method	$\Delta_{\text{Diameter}}$ [mm] ↓	$\Delta_{\text{CoM}}$ [mm] ↓	F1 Score [%] ↑
<b>Poke</b>	$6.3 \pm 0.6$	$6.0 \pm 0.5$	$45.8 \pm 0.6$
<b>Pri</b>	$2.8 \pm 0.3$	$2.3 \pm 0.1$	$74.7 \pm 1.2$
<b>Pri + Poke</b>	$2.0 \pm 0.2$	$1.6 \pm 0.1$	$80.9 \pm 0.8$

**Hand-Held Data** We finally tested LESS on data-handheld. To improve the results visually, we performed a calibration of the threshold for classifying a pixel in the image as an inclusion or a supporting pillar, choosing a threshold of 0.8. We tested LESS trained with data-primitive and LESS trained without data-primitive. As can be seen in Table IV, training on the combination of data-primitive and data-poke, helped mitigate the distribution shift when considering hand-held motion and resulted in improved performance. We achieve a  $\Delta_{\text{CoM}}$  of  $8.4[\text{mm}]$  compared to  $1.6[\text{mm}]$  in the robotic setup, about half a centimeter worse. Although our results show that tactile imaging using the robot-held sensor outperforms the hand-held setup, we believe that the results also show a positive trend: with more diverse robotic data, we expect to improve hand-held tactile imaging in the future.

TABLE IV: Testing LESS on data-handheld.

Method	$\Delta_{\text{Diameter}}$ [mm] ↓	$\Delta_{\text{CoM}}$ [mm] ↓	F1 Score [%] ↑
<b>w/ Pri,</b>	$2.8 \pm 0.8$	$8.4 \pm 0.7$	$41.3 \pm 2.9$
<b>w/o Pri,</b>	$5.2 \pm 0.5$	$9.5 \pm 0.8$	$34.7 \pm 4.2$

**Real-Time Hand-Held Tactile Imaging** We provide a visualization of our proposed hand-held tactile imaging system operation in Figure 10. A full demonstration video can be found on the project website ([zoharri.github.io/LESS](http://zoharri.github.io/LESS)).

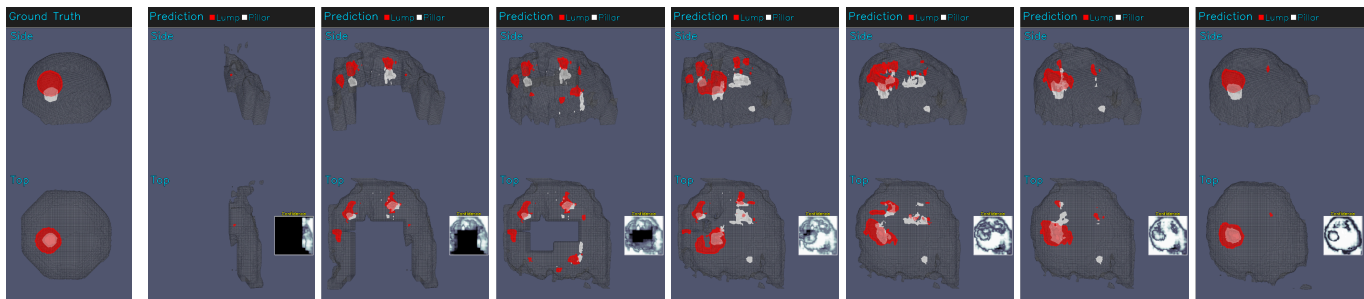


Fig. 10: The monitor presented to the operator at different steps while using our hand-held system. On the far left, we show the ground truth. The operator can observe the prediction and the uncertainty estimates (bottom right box in each plot).

## VI. DISCUSSION

Understanding artificial touch is a long-standing challenge in AI, with tactile imaging being an important application. In this work we take a step in this direction by proposing a neural network architecture tailored for spatio-temporal tactile data and new data collection techniques for making sense of soft objects. Technically, our work paves the way for learning-based tactile imaging that is hand-held, operator-independent, and applicable to arbitrary-shaped objects with composite internal structure. These capabilities have direct clinical value: operator independence is important for home-use devices, while detecting multiple inclusions can inform a physician’s decision to perform multiple biopsies, potentially improving patient outcomes.

A pressing question is clinical relevance of a method trained on synthetic data. In principle, with a large and diverse synthetic object dataset, learning-based tactile imaging may generalize to human tissue with sufficient accuracy. Future work must evaluate this hypothesis on human-subject data; our technical contribution in this work, in particular, the hand-held system we developed, will be important for clinical evaluation.

Beyond tactile imaging, our neural architecture may be useful for other robotic domains that involve interpreting complex spatio-temporal signals, such as learning object affordances from video [30], or soft object manipulation [26].

## ACKNOWLEDGMENTS

This work received funding from the European Union (ERC, Bayes-RL, Project Number 101041250). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

## REFERENCES

- [1] Iris Andrussov, Huanbo Sun, Katherine J Kuchenbecker, and Georg Martius. Minsight: A fingertip-sized vision-based tactile sensor for robotic manipulation. *Advanced Intelligent Systems*, 5(8):2300042, 2023.
- [2] Leonid Babarinow. Lift table, 2021. URL <https://www.printables.com/model/103257-lift-table-130x100>.
- [3] Roberto Calandra, Andrew Owens, Manu Upadhyaya, Wenzhen Yuan, Justin Lin, Edward H. Adelson, and Sergey Levine. The feeling of success: Does touch sensing help predict grasp outcomes? In *Conference on Robot Learning*, 2017.
- [4] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [5] Mauro Comi, Alessio Tonioni, Jonathan Tremblay, Max Yang, Valts Blukis, Yijiong Lin, Nathan F Lepora, and Laurence Aitchison. Snap-it, tap-it, splat-it: Tactile-informed 3d gaussian splatting for reconstructing challenging surfaces. In *2025 International Conference on 3D Vision (3DV)*, pages 1134–1143. IEEE, 2025.
- [6] Tal Daniel and Aviv Tamar. DDLP: Unsupervised object-centric video prediction with deep dynamic latent particles. *Transactions on Machine Learning Research (TMLR)*, 2023.
- [7] Daniel, Tal and Aviv Tamar. Unsupervised image representation learning with deep latent particles. In *International Conference on Machine Learning (ICML)*, 2022.
- [8] Chengjin Du, Federico Bernabei, Matteo Lo Preti, and Lucia Beccai. Design and optimization of a soft magnetic tactile sensor. In *2023 IEEE SENSORS*, pages 1–4. IEEE, 2023.
- [9] Vladimir Egorov and Armen P. Sarvazyan. Mechanical imaging of the breast. *IEEE Transactions on Medical Imaging*, 27(9):1275–1287, 2008.
- [10] Vladimir Egorov, Suren Ayrapetyan, and Armen P. Sarvazyan. Prostate mechanical imaging: 3-d image composition and feature calculations. *IEEE Transactions on*

- Medical Imaging*, 25(10):1329–1340, 2006.
- [11] Vladimir Egorov, Heather van Raalte, and Armen P. Sarvazyan. Vaginal tactile imaging. *IEEE Transactions on Biomedical Engineering*, 57(7):1736–1744, 2010.
- [12] Jean Elsner. Taming the panda with python: A powerful duo for seamless robotics programming and integration. *SoftwareX*, 24:101532, 2023. ISSN 2352-7110. doi: <https://doi.org/10.1016/j.softx.2023.101532>. URL <https://www.sciencedirect.com/science/article/pii/S2352711023002285>.
- [13] Ruoxuan Feng, Jiangyu Hu, Wenke Xia, Ao Shen, Yuhao Sun, Bin Fang, Di Hu, et al. Anytouch: Learning unified static-dynamic representation across multiple visuo-tactile sensors. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [14] Yang Gao, Lisa Anne Hendricks, Katherine J. Kuchenbecker, and Trevor Darrell. Deep learning for tactile understanding from visual and haptic data. In *International Conference on Robotics and Automation*, 2016.
- [15] Irmak Guzey, Ben Evans, Soumith Chintala, and Lerrel Pinto. Dexterity from touch: Self-supervised pre-training of tactile representations with robotic play. In *Conference on Robot Learning*, pages 3142–3166. PMLR, 2023.
- [16] Rory Hampson, Robert G. Anderson, and Gordon Dobie. Tactile imaging: The requirements to transition from screening to diagnosis of breast cancer—a concise review of current capabilities and strategic direction. Preprint, University of Strathclyde, 2022.
- [17] Carolina Higuera, Akash Sharma, Chaithanya Krishna Bodduluri, Taosha Fan, Patrick Lancaster, Mrinal Kalakrishnan, Michael Kaess, Byron Boots, Mike Lambeta, Tingfan Wu, et al. Sparsh: Self-supervised touch representations for vision-based tactile sensing. In *Conference on Robot Learning*, 2024.
- [18] George P Jenkinson, Karl Tiemann, Angeliki Papatthanasidou, Jonny Bewley, Andrew T Conn, and Antonia Tzemanaki. A robotic radial palpation mechanism for breast examination (iris). In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 337–342. IEEE, 2023.
- [19] Raghav Khanna and Elizabeth Shah. Robotics in screening, diagnosis and treatment of breast cancer: a perspective view. *Clinical Breast Cancer*, 24(1):17–26, 2024.
- [20] Tobias Kunz and Mike Stilman. Time-optimal trajectory generation for path following with bounded acceleration and velocity. In *Robotics: Science and Systems*, 2012. URL <https://api.semanticscholar.org/CorpusID:15630193>.
- [21] Mike Lambeta, Po-Wei Chou, Stephen Tian, Brian Yang, Benjamin Maloon, Victoria Rose Most, Dave Stroud, Raymond Santos, Ahmad Byagowi, Gregg Kammerer, et al. Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *IEEE Robotics and Automation Letters*, 5(3):3838–3845, 2020.
- [22] Sunghoon Lee, Sae Franklin, Faezeh Arab Hassani, Tomoyuki Yokota, Md Osman Goni Nayeem, Yan Wang, Raz Leib, Gordon Cheng, David W Franklin, and Takao Someya. Nanomesh pressure sensor for monitoring finger manipulation without sensory interference. *Science*, 370(6519):966–970, 2020.
- [23] Nathan F. Lepora, Alex Church, Conrad De Kerckhove, Raia Hadsell, and John Lloyd. From pixels to percepts: Highly robust edge perception and contour following using deep learning and an optical biomimetic tactile sensor. *IEEE Robotics and Automation Letters*, 4(2): 2101–2107, 2019. doi: 10.1109/LRA.2019.2899192.
- [24] Jianhua Li, Siyuan Dong, and Edward H. Adelson. Slip detection with combined tactile and visual information. In *International Conference on Robotics and Automation*, 2018.
- [25] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, pages 2999–3007, 2017. doi: 10.1109/ICCV.2017.324. URL <https://doi.org/10.1109/ICCV.2017.324>.
- [26] Xingyu Lin, Yufei Wang, Jake Olkin, and David Held. Softgym: Benchmarking deep reinforcement learning for deformable object manipulation. In *Conference on Robot Learning*, pages 432–448. PMLR, 2021.
- [27] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. *Advances in neural information processing systems*, 33:11525–11538, 2020.
- [28] Shan Luo, Wenzhen Yuan, Edward Adelson, Anthony G. Cohn, and Raul Fuentes. Vitac: Feature sharing between vision and tactile sensing for cloth texture recognition. In *ICRA*, 2018.
- [29] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 565–571, 2016. doi: 10.1109/3DV.2016.79. URL <https://doi.org/10.1109/3DV.2016.79>.
- [30] Tushar Nagarajan, Yanghao Li, Christoph Feichtenhofer, and Kristen Grauman. Ego-topo: Environment affordances from egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 163–172, 2020.
- [31] European Society of Radiology (ESR) communications@myESR.org Clevert Dirk-André Nyhsen Christiane Ricci Paolo Sidhu Paul S. Tziakouri Chrysa Radziņa Maija Cantisani Vito Brady Adrian P. Position statement and best practice recommendations on the imaging use of ultrasound from the european society of radiology ultrasound subcommittee. *Insights into Imaging*, 11(1): 115, 2020.
- [32] Haozhi Qi, Brent Yi, Sudharshan Suresh, Mike Lambeta, Yi Ma, Roberto Calandra, and Jitendra Malik. General in-hand object rotation with vision and touch. In *Conference*

- on *Robot Learning*, pages 2549–2564. PMLR, 2023.
- [33] Zohar Rimon and Aviv Tamar. Artificial palpation poke dataset. Zenodo Dataset, 2025. URL <https://doi.org/10.5281/zenodo.17608184>. Version v1; Automatic robotic-poking tactile dataset used in “Towards Artificial Palpation: Representation Learning of Touch on Soft Bodies” (NeurIPS 2025).
- [34] Zohar Rimon, Elisei Shafer, Tal Tepper, Efrat Shimron, and Aviv Tamar. Toward Artificial Palpation: Representation Learning of Touch on Soft Bodies. In *Advances in Neural Information Processing Systems (NeurIPS)*, December 2025.
- [35] Rimon, Zohar, Aviv Tamar, and Gilad Adler. Meta reinforcement learning with finite training tasks – a density estimation approach. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [36] Oluwatoyin Sanni, Giorgio Bonvicini, Muhammad Arshad Khan, Pablo C López-Custodio, Kiyanoush Nazari, et al. Deep movement primitives: toward breast cancer examination robot. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 12126–12134, 2022.
- [37] Armen Sarvazyan, Timothy J Hall, Matthew W Urban, Mostafa Fatemi, Salavat R Aglyamov, and Brian S Garra. An overview of elastography—an emerging branch of medical imaging. *Current Medical Imaging*, 7(4):255–282, 2011.
- [38] Armen P. Sarvazyan and Vladimir Egorov. Mechanical imaging—a technology for 3-d visualization and characterization of soft tissue abnormalities: A review. *Current Medical Imaging Reviews*, 8(1):64–73, 2012.
- [39] Luca Scimeca, Josie Hughes, Perla Maiolino, Liang He, Thrishantha Nanayakkara, and Fumiya Iida. Action augmentation of tactile perception for soft-body palpation. *Soft robotics*, 9(2):280–292, 2022.
- [40] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun Woo. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *NeurIPS*, 2015.
- [41] Sudharshan Suresh, Haozhi Qi, Tingfan Wu, Taosha Fan, Luis Pineda, Mike Lambeta, Jitendra Malik, Mrinal Kalakrishnan, Roberto Calandra, Michael Kaess, et al. Neuralfeels with neural fields: Visuotactile perception for in-hand manipulation. *Science Robotics*, 9(96):ead10628, 2024.
- [42] Togzhan Syrymova, Amir Yelenov, Karina Burunchina, Nazgul Abulkhanova, Huseyin Atakan Varol, Juan Antonio Corrales Ramon, and Zhanat Kappasov. Breast lump detection and localization with a tactile glove using deep learning. *arXiv preprint arXiv:2502.15767*, 2025.
- [43] Tito Pradhono Tomo, Massimo Regoli, Alexander Schmitz, Lorenzo Natale, Harris Kristanto, Sophon Somlor, Lorenzo Jamone, Giorgio Metta, and Shigeki Sugano. A new silicone structure for uSkin—A soft, distributed, digital 3-axis skin sensor and its integration on the humanoid robot iCub. *IEEE Robotics and Automation Letters*, 3(3):2584–2591, 2018.
- [44] Robert E. Weiss, Vladimir Egorov, Suren Ayrapetyan, Noun Sarvazyan, and Armen Sarvazyan. Prostate mechanical imaging: A new method for prostate assessment. *Urology*, 71(3):425–429, 2008.
- [45] Peter S. Wellman, Elizabeth P. Dalton, David Krag, Kenneth A. Kern, and Robert D. Howe. Tactile imaging of breast masses: First clinical report. *Archives of Surgery*, 136(2):204–208, 2001.
- [46] Wenqiang Xu, Zhenjun Yu, Han Xue, Ruolin Ye, Siqiong Yao, and Cewu Lu. Visual-tactile sensing for in-hand object reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8803–8812, 2023.
- [47] Wenzhen Yuan, Yuchen Mo, Shaoxiong Wang, and Edward H. Adelson. Active clothing material perception using tactile sensing and deep learning. In *ICRA*, 2018.

APPENDIX

A	Phantom Fabrication Technical Details . . . . .	12
B	Ground-Truth Preparation and Alignment . . . . .	12
C	Tactile Imaging Metrics . . . . .	12
D	Imaging Class Imbalance . . . . .	15
E	LESS Self-Supervised Loss . . . . .	16
F	Pose Estimation System Technical Details . . . . .	16
G	Pose Estimation Accuracy . . . . .	17
H	Automatic Data Collection . . . . .	17
I	Human Primitives Based Data Collection . . . . .	18
J	Hand-Held Technical Details . . . . .	18

A. Phantom Fabrication Technical Details

In [34], silicone was spread on the mold in much the same fashion as in the manufacturing of hollow chocolates. We chose to use a different process where a consistent, uniform thickness for the skin of the phantom is achieved by using two molds. We will describe the differences in the manufacture of each part in the following sections.

**Shell** In the [34] process, silicone was added to a mold and put into a tumbler for an even coat. In our process we have two separate molds, one for the base with a cavity for the insert (see Figure 16a), another for the top skin (see Figure 16b). Note that the cavity is molded into the base plate. Silicone is poured into each of these molds in parallel. After the silicone has set, we do the following with each of the molds: For the top mold we remove only the top part while the silicone is still attached to the bottom part. For the cavity mold, we remove the silicone cavity while leaving it attached to the base plate as in Figure 16a. We then pour silicone onto the base plate to achieve a thickness of a few millimeters and cover with the bottom part still attached to the mold, as in Figure 16c

**Insert** As with the shell, we have separate molds, in this case one for the top part and another for the inclusion (see Figures 15a and 15b). The two-part mold for the insert also has the added benefit of enabling the manufacture of more complex inclusions (e.g. Figure 14).

To manufacture the insert, we pour silicone into the two-part mold to form the top skin. In parallel, we pour silicone into the inclusion’s mold. We let the silicone cure and, as with the shell, we remove the inner part of the mold but leave the cured silicone attached to the outer part, as in Figure 15a. We remove the inclusion from its mold and put it into a designated hole in the base. Then, as with the shell, we pour silicone onto the base so that it is a few millimeters high, let it seep into the anchor holes, and cover it with the top part with silicone attached as in Figure 15c.

We have manufactured the inserts from [34] as well as new ones (with similar inclusions), all of which are shown in Figure 11. In addition, we manufacture inserts to test the generalization of LESS all of which are shown in Figure 12.

**Lift** We tried collecting data in the same fashion as in [34], but found the process to be slow, since after collecting data for a single configuration we would need to remove the insert

manually and put it in again. Since the base is an octagon, we can collect eight different configurations for one insert. We decided to create a mechanism to lower the insert from under the shell, rotate it, and then put it back in place (see Figure 13).

The mechanism created is a motorized scissor lift based on print files for a manual lift [2]. We added a Dynamixel servo to drive the lift and another to rotate the insert. To further streamline data collection, we use 3 lifts for a single robot arm, enabling data collection without human intervention for hours on end.

B. Ground-Truth Preparation and Alignment

As discussed in Section IV-C, each pixel in our imaging task is assigned to one of four classes: background, insert, pillar, or inclusion. To generate our ground truth, we developed an MRI processing pipeline, along with a visualization tool to interpret the ground truth and the predictions. An overview of this procedure is shown in Figure 18.

C. Tactile Imaging Metrics

Whereas prior work by Rimon et al. [34] demonstrated a proof of concept for artificial palpation using 2D tactile images, it has a key limitation: the resulting representations and tactile images are restricted to two dimensions and do not capture the 3D structure of objects, limiting information about the size, shape, and position of internal inclusions as shown in Figure 19 (a). We address this limitation by generating 3D tactile images, enabling the tactile data to capture a richer and more precise representation of the objects.

In Section V-A, we compare predictions obtained from a full 3D reconstruction with predictions evaluated in a purely 2D setting. We report results for both 2D predictions and 3D predictions, while ensuring that evaluation is conducted under comparable conditions. In particular, metrics are computed either directly on the full 3D volumetric prediction or on a single 2D slice corresponding to the slice with the largest inclusion area. The latter follows the evaluation introduced in [34] and enables a fair and direct comparison between 2D and 3D approaches under identical slice-level conditions.

We focus our evaluation on the inclusion class and define the following metrics, which are computed via element-wise comparison at the voxel level (for 3D volumes) and pixel level (for 2D slices):

- **F1 score** - The F1 score is defined as the harmonic mean of precision and recall,

$$F1 = \frac{2 \cdot \text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$$

- **Center of mass (COM) error** - The COM error is defined as the Euclidean distance between the predicted inclusion center of mass and the ground truth inclusion center of mass.
- **Absolute Relative Size Error (ARSE)** - Let  $|p|$  denote the number of pixels classified as inclusion in the prediction and  $|y|$  the number of inclusion pixels in the ground




Fig. 11: All our manufactured single-inclusion inserts.




Fig. 12: All our manufactured multi-inclusion inserts.





Fig. 13: Scissor lift. A motor drives a screw to lift the contraption. On the top black platform, there is another motor to rotate the insert.



(a) Molds for top skin and skin.



(b) Mold for inclusion and ready inclusion





Fig. 14: Two-part mold for three inclusions in a curved formation




(c) Final Step

Fig. 15: Steps of the insert manufacturing process. We create the top skin and insert in parallel (a and b). The final step of the manufacturing process insert is shown in (c). The top part has the silicone skin inside it. For the bottom part, we put in the insert, add a flange (in dark grey) and pour silicone (in transparent light grey).

truth. The Absolute Relative Size Error is defined as

$$ARSE = \frac{||p| - |y||}{\frac{1}{2} (|p| + |y|)}$$

- **Absolute Diameter Error** - We define the diameter of an inclusion as the maximum Euclidean distance, measured



(c) Shell mold final step

Fig. 16: Manufacturing process for the shell: We cast the cavity (a) and outer skin (b) in parallel (cured silicone in light grey). After curing, we pour silicone on the base plate and join it with the outer skin mold with cured silicone (c).

in millimeters, between any two pixels classified as an inclusion within the same slice. The diameter is computed for each slice independently, and the final inclusion diameter is defined as the maximum diameter over all slices. The Absolute Diameter Error is then defined as the absolute difference between the predicted diameter and the corresponding ground-truth diameter.

In Table I, we observe that predicting the full 3D structure not only improves the interpretability of the results for human observers, but also leads to consistent improvements across all 2D metrics. These findings suggest that the added geometric information captured by the 3D reconstruction provides measurable benefits beyond visualization alone.




Fig. 17: On the right, a regular-sized phantom. On the left, a large phantom, which is a  $\times 2$  scaled version of the regular phantom in the x and y directions.

#### D. Imaging Class Imbalance

In [34], a cross-entropy loss for all the image pixels was used. This loss can be directly extended to voxel predictions. However, we found that when extending the formulation to 3D reconstruction, in particular, moving from 2D slices introduced a substantial increase in class imbalance, as the proportion of voxels belonging to the inclusion class decreased significantly relative to the background and phantom classes. This shift is quantified in Table V, which reports the mean and standard deviation of the number of pixels (or voxels) per class across the dataset.

Class	2D (%)	3D (%)
Background	$47.40 \pm 2.04$	$50.97 \pm 2.17$
Phantom	$49.78 \pm 2.07$	$47.71 \pm 2.17$
Pillar	—	$0.33 \pm 0.04$
Inclusion	$2.82 \pm 1.01$	$0.99 \pm 0.48$

TABLE V: Per-Pixel Class Distribution

In the presence of severe class imbalance, cross-entropy loss has been shown to struggle [25], resulting in insufficient learning performance for minority classes. This observation motivated us to adopt a loss function that directly addresses class imbalance. In particular, we use a linear interpolation between **focal loss** and **Dice loss**, which combines the strengths of both approaches. Focal loss is designed to tackle extreme class imbalance by down-weighting easy-to-classify examples, while Dice loss optimizes the overlap between different classes, which encourages more precise boundaries.

We will now formally define our loss function. To ensure balanced gradients during optimization, we employ a dynamic scaling strategy. The total loss is defined as:

$$\begin{aligned} \mathcal{L} &= \lambda \cdot \mathcal{L}_{\text{Dice}} + (1 - \lambda) \cdot \eta \cdot \mathcal{L}_{\text{Focal}}, \\ \eta &= \frac{\mathcal{L}_{\text{Dice}}}{\mathcal{L}_{\text{Focal}} + \delta}, \\ \mathcal{L}_{\text{Focal}} &= -\frac{1}{HWD} \sum_{i,j,k} \sum_{c=1}^C \alpha_c y_{ijk,c} (1 - p_{ijk,c})^\gamma \log(p_{ijk,c}), \\ \mathcal{L}_{\text{Dice}} &= 1 - \frac{1}{C} \sum_{c=1}^C \frac{2 \sum_{i,j,k} p_{ijk,c} y_{ijk,c} + \epsilon}{\sum_{i,j,k} p_{ijk,c} + \sum_{i,j,k} y_{ijk,c} + \epsilon} \end{aligned}$$

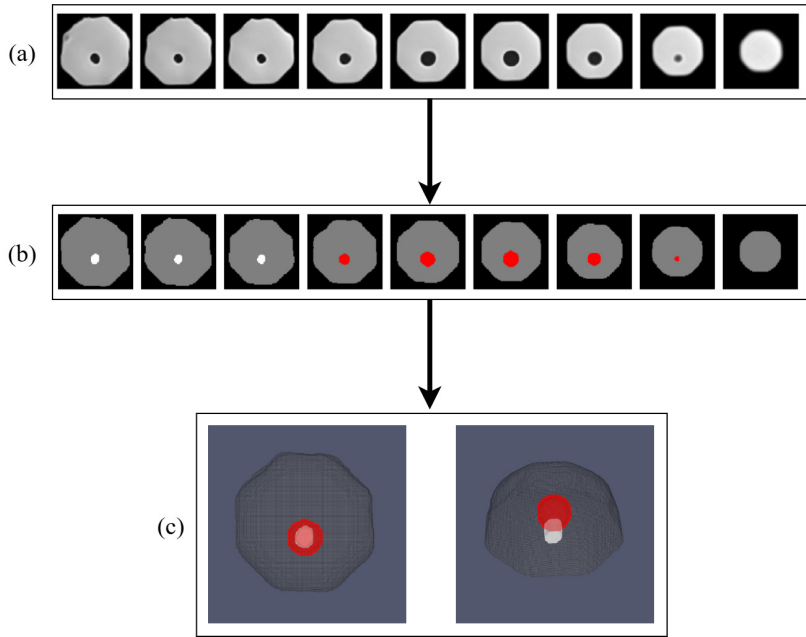


Fig. 18: Overview of our pre-processing pipeline: (a) A subset of slices from the original MRI scan of a single phantom. (b) The corresponding processed scan with pixel-wise class annotations, where black denotes background, gray denotes the phantom insert, white denotes the pillar, and red denotes the inclusion. (c) The processed MRI scan is then sent to our custom-made visualization program, which renders the 3D scan.




Fig. 19: Comparison between 2D and 3D visualization approaches. (a) Representative 2D slice-based visualization as presented in [34]. (b) Top-view rendering of the proposed 3D visualization. (c) Side-view rendering of the same 3D visualization, highlighting the vertical connectivity between the pillar and the inclusion.

Here,  $H$ ,  $W$ , and  $D$  denote the height, width, and depth of the input volume, respectively, and  $C$  represents the number of classes (4 in our case). The variable  $y_{ijk,c}$  denotes the one-hot encoded ground truth (equal to 1 if the voxel at  $(i, j, k)$  belongs to class  $c$ , and 0 otherwise), while  $p_{ijk,c}$  represents the corresponding predicted softmax probability. To prevent the Dice loss from dominating the optimization, we introduce a scaling factor  $\eta$ , which dynamically normalizes the Focal loss magnitude to match the Dice loss at each iteration. We set  $\lambda = 0.5$  and the focusing parameter  $\gamma = 2$ . The class-balancing weights  $\alpha_c$  are calculated based on the inverse frequency of each class within the batch. Finally,  $\epsilon = 10^{-6}$  and  $\delta = 10^{-8}$  are included for numerical stability.

#### E. LESS Self-Supervised Loss

We adapt the force-prediction self-supervised loss from Rimon et al. [34] to our local prediction.

To allow for parallel computation of all particles' sequences, after filtering the sequences according to the receptive fields, we pad each local input sequence with the last measurement to a constant length  $L$ . In the following reconstruction procedure, we do not sample steps involving padding measurements.

We use a mean squared error (MSE) reconstruction loss between the predicted and true forces. We uniformly subsample reconstruction steps as described in [34] to mitigate memory complexity:

$$\mathcal{L}_{rec} = \frac{1}{2NKK'} \sum_{n=1}^N \sum_{k=1}^K \sum_{k'=1}^{K'} \left\| FD \left( z_{t_k^n}^{\bar{x}_n}, x_{t_{k'}^n} - \bar{x}_n \right) - f_{t_{k'}^n} \right\|^2,$$

where FD is the Force Decoder, as described in Section IV-A, for each  $n \in [1, N]$ ,  $\{t_k^n\}_{k=1}^K$  are  $K < L$  uniform samples without replacement from the non-padded input measurements of  $\bar{x}_n$ , and for each  $k \in [1, K]$ ,  $\{t_{k'}^n\}_{k'=1}^{K'}$  are also  $K' < L$  uniform samples. We set  $K = K' = 16$ .

#### F. Pose Estimation System Technical Details

Our setup comprises a tactile sensor (Xela uSkin) mounted on a custom 3D-printed, hand-held device. During operation, the user palpates the phantom using the sensor to acquire tactile data, while the model generates a real-time heat map localizing hard inclusions within the phantom. Accurately

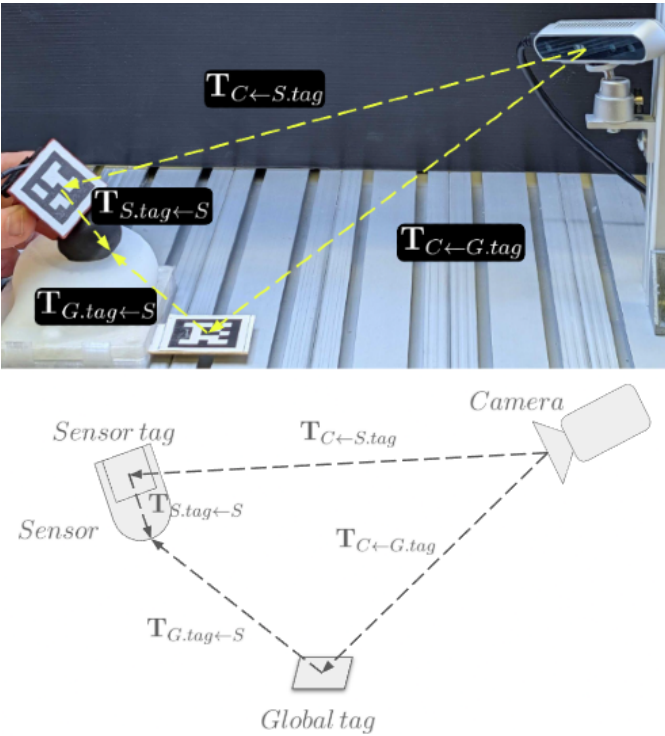


Fig. 20: Pose estimation system

measuring each sample’s spatial position of the sensor is required for our model.

To this end, we employ a visual tracking system using two AprilTags (each  $31 \times 31$  mm): a *sensor tag* mounted on the hand-held device and a *global tag* placed near the phantom as a static reference. These tags are continuously monitored by a calibrated, fixed Intel RealSense 415 camera as shown in Figure 20. Based on empirical optimization, the camera is configured to operate at 30 FPS, with a fixed exposure time of 4 ms and a sensor gain of 160 (device units), utilizing the left infrared (IR) imager (emitter disabled) for robust tag detection.

To estimate the spatial pose of the sensor relative to the phantom, we use a chain of rigid body transformations in  $SE(3)$ .

Let  $\mathbf{T}_{A \leftarrow B} \in SE(3)$  be a rigid transform that maps points expressed in frame  $B$  into frame  $A$  (i.e.,  $\mathbf{x}_A = \mathbf{T}_{A \leftarrow B} \mathbf{x}_B$  in homogeneous coordinates). We denote the transformation from the sensor tag to the sensor as  $\mathbf{T}_{S.tag \leftarrow S}$  and obtain it via the CAD model of the device. During sampling, the camera estimates the transformations to both the sensor and global tag denoted as  $\mathbf{T}_{C \leftarrow S.tag}$  and  $\mathbf{T}_{C \leftarrow G.tag}$ . By the identity  $\mathbf{T}_{C \leftarrow G.tag}^{-1} = \mathbf{T}_{G.tag \leftarrow C}$  we calculate the sensor’s pose in the global reference frame as:

$$\mathbf{T}_{G.tag \leftarrow S} = \mathbf{T}_{G.tag \leftarrow C} \cdot \mathbf{T}_{C \leftarrow S.tag} \cdot \mathbf{T}_{S.tag \leftarrow S}.$$

To calibrate the pose of the sensor from the cameras to the robot base frame, as illustrated in Figure 21, we used a standard Kabsch algorithm to find a transformation  $X$  such that  $X \cdot A^{-1} \cdot B \cdot C = D$ .  $A, B$  are obtained from April tags (using pupil-apriltags; camera calibration parameters obtained

by realsense SDK),  $C$  from the CAD of our 3D-printed rig, and  $D$  from robot kinematics.

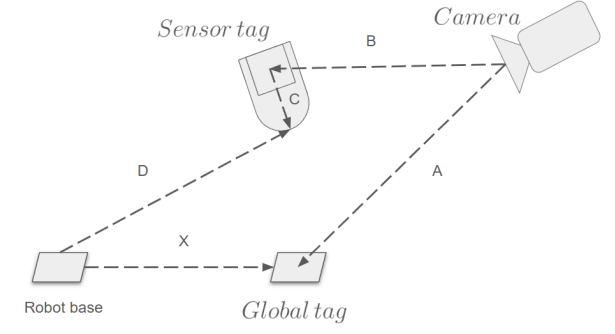


Fig. 21: Calibration procedure of pose estimation to the robot’s system.

### G. Pose Estimation Accuracy

To evaluate the accuracy of the system, we recorded the estimated pose during the sampling sequence and compared it against the pose from the robot’s kinematics, which serves as the ground truth. This yielded two time-indexed trajectories representing the device’s motion.

Given that the robot’s reporting frequency exceeds that of our optical system, the trajectories were temporally synchronized via interpolation. Assuming the error follows a zero mean independent Gaussian distribution, we computed the empirical  $\sigma$  as:  $\hat{\sigma}_{positional} = \sqrt{\frac{1}{3N} \sum_{i=1}^N (e_{i,x}^2 + e_{i,y}^2 + e_{i,z}^2)}$  for both positional and angular errors. The results show  $\hat{\sigma}_{positional} \leq 0.45\text{mm}$  and  $\hat{\sigma}_{angular} \leq 0.012\text{ radians}$ . These satisfy the accuracy requirements established in the previous section.

As the system’s error is unlikely to follow an independent Gaussian distribution in practice, we opted for an end-to-end evaluation of performance degradation. We ran the model on a sample collected using both the pose estimation and the robot’s true pose (as above) and evaluated the LESS F1 score under each condition. The model achieved F1 scores of  $67.6 \pm 0.0$  and  $67.3 \pm 0.1$  when using the robot’s true pose and the estimated pose, respectively.

### H. Automatic Data Collection

We have performed our data collection using a Franka Emika Panda robot and a slightly modified panda\_py package [12]. In order to collect a vast and diverse dataset, we implemented an automated sampling system. Insert rotation was automated using custom 3D-printed lifts equipped with rotating ramps as explained in Appendix A. Pairing this capability with adding randomization to the sampling parameters facilitated continuous data acquisition 24/7 since we had 3 elevators, so we had sweeps of about 8 hours on 3 phantoms, and because each scan is different from the other, collecting more data on the same phantoms is still beneficial, increasing the variance of the resulting dataset. The randomness was introduced by applying spatial noise to the grid coordinates (smaller than

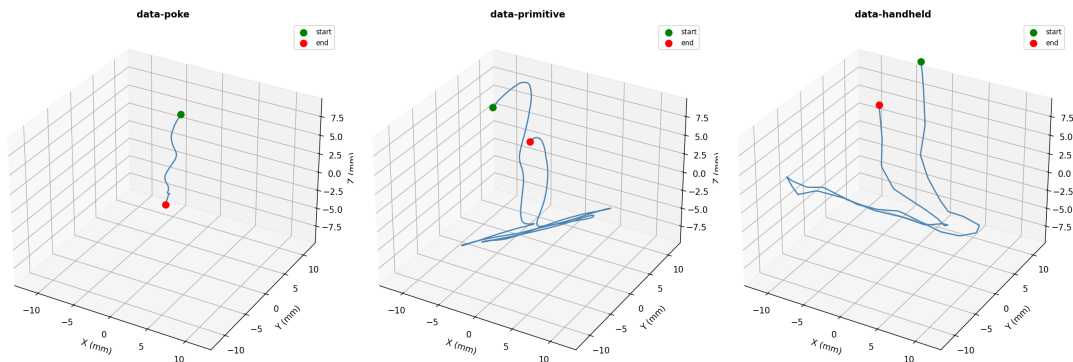


Fig. 22: Comparison of a sampled trajectory from the different datasets. The primitives are more similar to the data-handheld, compared to data-poke

the grid step) and stochastic adjustments to the end effector yaw angle (0 – 80 degrees). We have collected 3372 scans with a random yaw angle and randomized positions. The scans have been performed on 8 shells, 54 inserts, and with 3 Xela sensors. The data collection was done with the same controller as in Rimon et al. [34].

### I. Human Primitives Based Data Collection

In order to train the model to be able to work well with a hand-held sensor, we needed a way to imitate human motions. Our solution was to record several human-generated primitive motions using the robot. An operator performed several basic palpation motions on the phantom while the robot’s poses were recorded. Those primitive trajectories were then used instead of the robot pokes. While sampling each grid point, we choose a random primitive and calculate a trajectory that will originate there. Each trajectory is evaluated to ensure all points remain within the predefined boundaries of the specific phantom. If a trajectory exceeds these limits but maintains sufficient length, it is truncated; otherwise, a different motion primitive is selected. Three basic primitives were used: a poke, a back-and-forth palpation movement, and a left-and-right palpation movement. Videos of these motions are available on the website. A sampled trajectory from the primitives, compared to the poke-data and poke-handheld can be seen in Figure 22. To run those trajectories on the robot, we used `panda_py` by Elsner [12]. We used the `move_to_joint_position` function, which calculates detailed movements for the robot from the trajectory using the TOTG algorithm by Kunz and Stilman [20]. We have sampled a total of 1409 full scans using the motion primitives method.

### J. Hand-Held Technical Details

**System Latency:** The model-forward pass is  $250ms$ ; other system latencies are negligible in comparison, allowing  $\sim 4[Hz]$  predictions for real time interaction

**Data Collection:** data-handheld contains 20 complete scans (approximately 2000 single trajectories).