

# Few-Shot Learning of Force-Based Motions From Demonstration Through Pre-training of Haptic Representation

Marina Y. Aoyama<sup>1</sup>, João Moura<sup>1,2</sup>, Namiko Saito<sup>1,2</sup> and Sethu Vijayakumar<sup>1,2</sup>

**Abstract**—In many contact-rich tasks, force sensing plays an essential role in adapting the motion to the physical properties of the manipulated object. To enable robots to capture the underlying distribution of object properties necessary for generalising learnt manipulation tasks to unseen objects, existing Learning from Demonstration (LfD) approaches require a large number of costly human demonstrations. Our proposed semi-supervised LfD approach decouples the learnt model into a haptic representation encoder and a motion generation decoder. This enables us to pre-train the first using a large amount of unsupervised data, easily accessible, while using few-shot LfD to train the second, leveraging the benefits of learning skills from humans. We validate the approach on the wiping task using sponges with different stiffness and surface friction. Our results demonstrate that pre-training significantly improves the ability of the LfD model to recognise physical properties and generate desired wiping motions for unseen sponges, outperforming the LfD method without pre-training. We validate the motion generated by our semi-supervised LfD model on the physical robot hardware using the KUKA iiwa robot arm. We also validate that the haptic representation encoder, pre-trained in simulation, captures the properties of real objects, explaining its contribution to improving the generalisation of the downstream task. See our accompanying video: <https://youtu.be/zP4JvHaCWHk>.

## I. INTRODUCTION

Humans are able to adapt their motion according to the different physical properties of the objects they manipulate by extensively using force and tactile sensing [1]–[3]. For instance, when wiping a surface with a sponge, humans naturally adapt their wiping motion based on the stiffness and friction of both the table and sponge, in order to maximise cleaning effectiveness while avoiding damage [4], [5]. Similarly, in various other robotic contact-rich manipulation tasks such as food manipulation [6]–[8], pouring [9], [10], or peg-in-hole insertion [11], the ability to adapt the motion according to the physical properties of the manipulated objects is vital for the successful execution of the task.

Learning from Demonstration (LfD) is an intuitive and efficient approach for robots to acquire skills or behaviours by imitating human demonstrations [12]–[14]. One of the challenges in LfD is how to generalise learnt manipulation motions to objects with different physical properties. In addition, producing human demonstrations requires some cognitive and physical effort, typically resulting in few examples with a small variety of objects [15], [16]. Therefore,

<sup>1</sup>Authors are with the School of Informatics, The University of Edinburgh, Edinburgh, U.K.

<sup>2</sup>Authors are with The Alan Turing Institute, London, U.K.

This work is supported by the EU H2020 Project Harmony (Grant No. 101017008), JST Moonshot R&D (Grant No. JPMJMS2031), the Kawada Robotics Corporation, and The Alan Turing Institute.

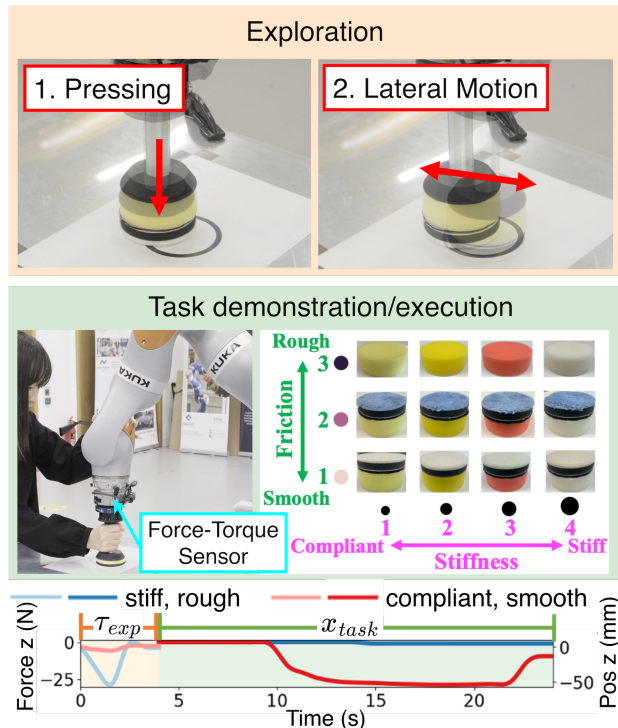


Fig. 1: Wiping Experiment. In the exploration phase, the robot performs exploratory actions to identify the physical properties of the sponge. Then, in the task execution phase, the robot performs the desired wiping motion according to the physical properties of the sponge identified in the exploration phase.

we address the question – how can robots learn to manipulate unseen objects with different physical properties from a small number of human demonstrations?

### A. Related Work

There are previous works in LfD that consider physical properties of objects in order to adapt the learnt motion [7], [10], [17]. These works achieve adaptation by performing exploratory actions on the manipulated object to obtain the latent state of the learnt target task, with the assumption that this latent state implicitly encodes the object characteristics. However, there is a coupling between the exploration and target tasks in the sense that the robot needs to jointly learn both the exploration and the task models, the first responsible for characteristic encoding and the second for motion generation. Therefore, in order to capture the underlying distribution of object properties and improve generalisation, one needs to increase the number of demonstrations, incurring a significant amount of human effort.

In the context of representation learning, there are multiple studies addressing the pre-training of perceptual models using unsupervised data [18], [19]. Particularly in robotics, pre-training visual representations leads to improved robustness of the learnt tasks to different image observations [20]–[22]. While visual representation learning benefits from widely available image datasets [23] or pre-trained models [24], what remains unclear is how to acquire large-scale unsupervised data that effectively captures the physical properties of the objects and how to pre-train models that encode these properties using force information.

Force or tactile perception can reveal physical properties of objects, such as weight, texture, and compliance that are typically unavailable through visual perception [25]–[28]. To address this, Guzey, Evans, Chintala, and Pinto [29] pre-train a representation based on tactile observations. However, none of these visual [20], [21] or tactile [29] approaches apply self-supervised learning on the time series of observations. In contrast, there are works that highlight the importance of force or tactile time-series, obtained through interactive exploration, in effectively capturing physical object properties [27], [30]–[32], albeit without ever using that information to adapt motion learnt from human demonstration. Therefore, we propose a haptic representation that considers a series of force observations and address the question of how to adapt motion based on the obtained object characteristics.

### B. Problem Formulation

Similarly to the prior works on adapting motion based on the physical properties of the object [9], [10], [17], we consider a task with two phases: the exploration phase, where the robot performs pre-defined exploratory actions to identify the properties of the object, and the task execution phase, where the robot performs the target task according to the properties of the object identified in the exploration phase, as shown in Fig. 1.

The objective of LfD is to enable robots to perform manipulation tasks similar to human demonstrators. The robot must learn to adapt target forces based on the physical properties of the object, making a force tracking controller, requiring knowledge of the target force profile, unsuitable. Therefore, we aim to learn a model  $\theta_{LfD}$  that, given the force trajectory  $\tau_{exp} \in \mathbb{R}^{T_{exp} \times E_{exp} \times D_{exp}}$  obtained through the exploration, generates a motion trajectory  $x_{task} \in \mathbb{R}^{T_{task} \times D_{task}}$  for the target task, similar to the human demonstration. Here,  $T_{exp}$ ,  $E_{exp}$ , and  $D_{exp}$  represent the time steps, exploratory actions, and sensor information dimension, respectively. Similarly,  $T_{task}$  and  $D_{task}$  denote the time steps for task execution and the action dimension, respectively.

### C. Contribution

To address the issue of the existing LfD approaches that suffer from limited generalisation with few demonstrations, we propose a semi-supervised LfD approach that decouples the haptic representation encoder – the model representing the physical properties of an object – from the motion generation decoder, which maps the encoded object properties

to the desired behaviour. Through pre-training of the haptic representation encoder enabled by this decoupling, the aim is to utilise a large amount of unsupervised haptic data for generalising the task to unseen objects with different physical properties, even for limited number of demonstrations, while leveraging the benefits of learning skills from humans through LfD.

The training data then consists of two types of data sets: unsupervised haptic data and demonstration data. The unsupervised haptic data consists of  $M$  force trajectories  $\tau_{exp}$  for objects with different properties, autonomously collected by repeatedly performing pre-defined exploratory actions either in simulation or on a real robot. Note that we exclude the labels of the properties from the data set. The demonstration data consists of  $N$  pairs of trajectories, one for each object, composed of a force trajectory  $\tau_{exp}$ , autonomously collected as above, and the corresponding motion trajectory  $x_{task}$ , demonstrated by an expert demonstrator to perform the target task.  $M$  is much larger than  $N$  – the number of available human demonstrations.

We validate the proposed semi-supervised LfD approach on a wiping task using sponges with different stiffness and surface friction. The task serves as an exemplar force-based manipulation task, as the desired wiping motion depends on the physical properties of the sponge perceived only through observing the interaction force. Our results suggest that being able to pre-train the haptic representation encoder improves the learning of the wiping motions for unseen sponges. Moreover, replaying these learnt motions on the physical robot results in force interactions closer to those demonstrated by humans. Furthermore, a t-distributed Stochastic Neighbor Embedding (t-SNE) analysis of the model pre-trained in simulation supports the assumption that its latent space indeed captures the physical properties of objects.

In summary, the contributions of this work are:

- To the best of our knowledge, we propose the first semi-supervised LfD framework for pre-training of physical properties from time-series of force information.
- We propose a decoupling of the haptic representation encoder from the motion generation decoder, enabling the use of a large amount of unsupervised data to improve the generalisation capability of few-shot LfD for objects unseen during demonstration.
- We validate that the latent space of the haptic representation encoder, pre-trained in simulation, is able to distinguish the physical properties of the real objects, supporting the assumption that the latent space of our model indeed captures the object properties.
- We validate our semi-supervised LfD approach on a physical robotic setup, showing that it generates motion that achieves a force profile close to those demonstrated by the human.

## II. METHOD

To utilise unsupervised haptic data, we propose a semi-supervised LfD approach by decoupling the training of the haptic representation encoder  $\phi$  where  $z = \phi(\tau_{exp})$  and the

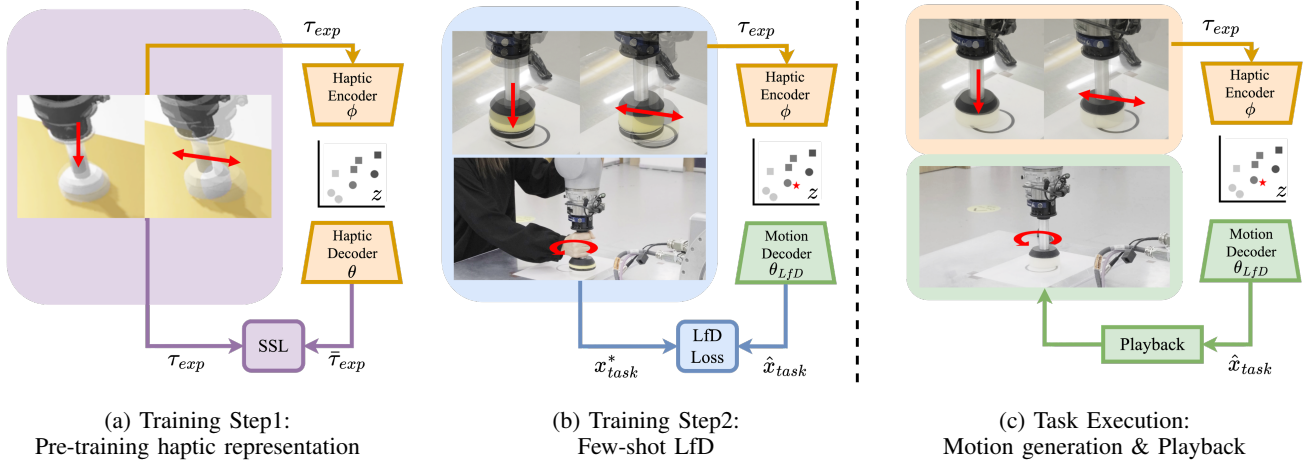


Fig. 2: Proposed semi-supervised LfD framework. First, we pre-train the haptic representation encoder on unsupervised haptic data collected in simulation by minimising the Self-Supervised Loss (SSL) Eq. (1). Then, we train the motion decoder to map from the haptic representation  $z$  obtained using the pre-trained haptic encoder to the desired motion for the target task on demonstration data by minimising the LfD loss Eq. (2).

motion generation decoder  $\theta_{LfD}$  where  $x_{task} = \theta_{LfD}(z)$ . In contrast to the approaches that require the joint training of the encoder  $\phi$  and the decoder  $\theta_{LfD}$ , the proposed decoupling enables training a part of the model, the haptic representation encoder  $\phi$ , using unsupervised haptic data. We expect the latent representation  $z$  to implicitly encode object properties. Note that the dimension of the latent space  $z$  is a hyperparameter and needs to be sufficiently large to be expressive yet small enough to prevent overfitting and to be computationally efficient. We explain the specific choice of the neural network hyperparameters for this work in Section III-D.

The proposed semi-supervised LfD framework consists of two steps: pre-training of the haptic representation encoder  $\phi$  using a large amount of unsupervised haptic data and a few-shot LfD of the motion generation decoder  $\theta_{LfD}$  using demonstration data, as illustrated in Fig. 2.

#### A. Pre-training haptic encoder

In the first step, we train a haptic representation encoder-decoder model. The objective of the model is to capture the underlying distribution of object properties in its latent space, rather than generating the desired motion for the target task, as shown in Fig. 2a. We adopt Variational Auto-Encoder (VAE) [33], a widely used approach for self-supervised learning of perceptual representation models, to train this haptic representation encoder-decoder model on unsupervised haptic data  $D_{offline} = \{\tau_{exp1}, \dots, \tau_{expM}\}$ . To achieve this, we reconstruct a force trajectory  $\tau_{exp}$  obtained through pre-defined exploratory actions and minimising a Self-Supervised Loss (SSL)

$$\hat{\theta}, \hat{\phi} = \arg \min_{\theta, \phi} L(\tau_{exp}) = \beta D_{KL}(q_{\phi}(z | \tau_{exp}) \| p_{\phi}(z)) + E_{MSE}(\tau_{exp}, \bar{\tau}_{exp}). \quad (1)$$

The loss function consists of two terms: the Kullback–Leibler (KL) divergence  $D_{KL}$  between the approximate posterior

and prior distribution of the latent variable  $z$ , and the reconstruction loss, which is the Mean Squared Error  $E_{MSE}$  between the input force trajectory  $\tau_{exp}$  and the reconstructed force trajectory  $\bar{\tau}_{exp}$ . Here,  $\phi$  and  $\theta$  represent the parameters of the haptic encoder and decoder, respectively, and  $\beta$  is the regularisation coefficient. Once we train the haptic representation encoder  $\phi$ , we freeze its weights and use them in the next step.

#### B. Few-shot LfD of motion decoder

In the second step, we train the motion generation decoder  $\theta_{LfD}$  to generate the desired motion for the target task  $x_{task}$  according to the properties of the manipulated object, using the demonstration data as depicted in Fig. 2b. For each demonstration in  $D_{demo} = \{(\tau_{exp1}, x_{task1}^*), \dots, (\tau_{expN}, x_{taskN}^*)\}$ , the haptic encoder  $\phi$ , which we pre-train in the first step, encodes the force trajectory  $\tau_{exp}$  obtained during the exploration phase into the haptic representation  $z$  of the object. Then, we train a mapping from the haptic representation  $z$  to the desired motion trajectory  $x_{task}$  for the target task by minimising the LfD loss

$$\hat{\theta}_{LfD} = \arg \min_{\theta_{LfD}} L(z, x_{task}) = E_{MSE}(\hat{x}_{task}, x_{task}^*) \quad (2)$$

where  $z \sim q_{\phi}(z | \tau_{exp})$ . Here,  $\theta_{LfD}$  represents the weights of the motion decoder. Since the objective is to generate motions for the target task that closely resemble those demonstrated for the object by a human demonstrator, we minimise the Mean Squared Error  $E_{MSE}$  between the generated motion trajectory  $\hat{x}_{task}$  and the demonstrated motion trajectory  $x_{task}^*$ .

#### C. Task Execution

At test time, the robot first performs the same pre-defined exploratory actions employed in training on the object to be manipulated before executing the task. The haptic encoder  $\phi$  encodes the force trajectory  $\tau_{exp}$  recorded during the

Stiffness											
Interpolation					Extrapolation						
		Stiffness level						Stiffness level			
		1	2	3	4			1	2	3	4
Friction level	1	o	x	x	o	Friction level	1	x	o	o	x
	2	o	x	x	o		2	x	o	o	x
	3	o	x	x	o		3	x	o	o	x

Surface friction											
Interpolation					Extrapolation						
		Stiffness level						Stiffness level			
		1	2	3	4			1	2	3	4
Friction level	1	o	o	o	o	Friction level	1	o	o	o	o
	2	x	x	x	x		2	o	o	o	o
	3	o	o	o	o		3	x	x	x	x

TABLE I: Training and test objects for stiffness interpolation, stiffness extrapolation, surface friction interpolation and surface friction extrapolation, with o: training and x: test.

exploration phase into a haptic representation  $z = \phi(\tau_{exp})$  of the object. Subsequently, the motion decoder  $\theta_{LfD}$  generates a motion trajectory  $x_{task} = \theta_{LfD}(z)$  specific to the object for the target task. Finally, the robot plays back the generated motion trajectory  $\hat{x}_{task}$  and executes the task without using real-time sensor feedback (*i.e.*, in an open-loop manner), as shown in Fig. 2c.

### III. EXPERIMENTAL SETUP

#### A. Wiping task

In our experiment, we use a wiping task as an instance of a force-based manipulation task as the desired wiping motion depends on the physical properties of the wiping tool.

The task is to wipe a flat table with a sponge attached to the end-effector of the robot arm, as shown in Fig. 1. We prepare sponges with different physical properties (*i.e.*, 4 levels of stiffness  $\times$  3 levels of surface friction) to evaluate the ability of the LfD model to adapt its wiping motion to the sponge properties unseen at training time.

#### B. System design and robot setup

We use a 7 Degrees of Freedom (DoF) KUKA iiwa robotic arm with a sponge attached to its end-effector for both simulation and real robot experiments. We mount a 6-axis force torque sensor on the wrist of the robot arm to measure force and torque along the x, y, and z axes. We control the robot by specifying the end-effector position when performing exploratory actions or replaying generated motions. We provide demonstrations by kinaesthetically moving the robot arm in gravity compensation mode. For the simulation, we use a physics simulation framework, ROS-PyBullet [34], [35], with exactly the same setup as the real robot experiment.

#### C. Data collection

1) *Unsupervised haptic data*: In the exploration phase, the robot sequentially performs two pre-defined exploratory actions: pressing (*i.e.*, moving the end-effector in the normal direction at a speed of 0.01 m/s for 2 seconds) and lateral motion (*i.e.*, moving the end-effector horizontally at speeds of 0.05 m/s and  $-0.05$  m/s for 1 second each). We selected these exploratory actions based on the intuition that they will yield meaningful information for identifying the physical

properties relevant to the wiping task. We record force and torque data along the x, y, and z directions at a frequency of 100 Hz during the exploration phase to obtain  $\tau_{exp} \in \mathbb{R}^{2400}$ .  $T_{exp}$ ,  $E_{exp}$ , and  $D_{exp}$  are 200, 2 and 6, respectively.

We collect 1000 unsupervised haptic data in simulation by performing the exploratory actions on simulated objects with varying contact stiffness  $k \in [80, 1000]$  N/m, lateral friction coefficient  $\mu_{lateral} \in [0.2, 8.0]$  and spinning friction coefficient  $\mu_{spin} \in [0.0, 4.0]$ . Note that a robot can obtain unsupervised haptic data either in simulation or on a physical robotic setup by repeatedly performing the pre-defined exploratory actions on a variety of objects. Data informativeness increases from unsupervised simulated data to unsupervised real data and demonstration data, with a corresponding increase in the cost of data collection. Unsupervised real data eliminates the issue of the sim2real gap; however, it is still more expensive to collect than unsupervised simulated data.

2) *Demonstration data*: First, the robot performs pre-defined exploratory actions on a sponge to obtain the exploration data  $\tau_{exp}$ , following the same procedure as in unsupervised haptic data collection. Subsequently, a human demonstrator kinaesthetically performs the desired wiping motion for the sponge in gravity compensation mode. The demonstrator wipes the table in a circular pattern, exerting as much force as possible in the normal direction to maximise cleaning effectiveness while maintaining a smooth motion. Therefore, the desired force depends on the stiffness and surface friction of the sponge, requiring less force for a stiff or rough sponge and vice versa. Each wiping motion lasts for 20 seconds. The position of the end-effector in the x, y, and z axes is recorded at a frequency of 100 Hz during the demonstration to obtain the desired wiping motion trajectory for the sponge, denoted as  $x_{task} \in \mathbb{R}^{6000}$ .  $T_{task}$  and  $D_{task}$  are 2000 and 3, respectively.

We collect demonstrations for all 12 sponges (4 stiffness levels  $\times$  3 surface friction levels  $\times$  1 trial), and left out the test cases for evaluation. We train the model on 4 different sets of training data ( $N = 6$  or 8), to test 4 different cases of interpolation and extrapolation of stiffness and surface friction, as shown in Table I.

#### D. Model training

Prior to the training, we apply a Butterworth low-pass filter to all force trajectories. The filter parameters used are as follows: sample rate  $r = 70000$ , cutoff frequency  $f_p = 1000$  Hz, stopband frequency  $f_s = 100$  Hz, maximum allowed loss in the passband  $g_{pass} = 3$  dB and maximum required loss in the stopband  $g_{stop} = 40$  dB. We normalise all force and motion trajectory data to  $[0, 0.9]$ .

First, we pre-train the haptic encoder, as described in Section II-A, using 1000 simulated unsupervised haptic data. The haptic encoder-decoder model consists of 1 fully connected encoder layer, 1 sampling step, and 1 fully connected decoder layer with Rectified Linear Unit (ReLU) as the activation function. To prevent overfitting, we apply a dropout rate of 0.1 to the decoder. We train the haptic encoder-

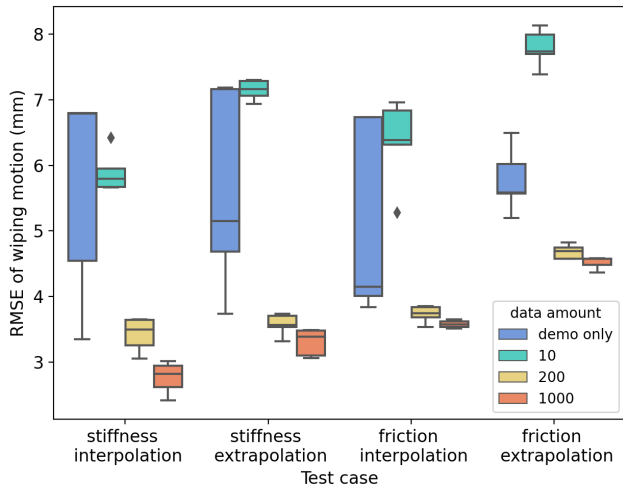


Fig. 3: RMSE between the wiping motion trajectories demonstrated and generated by the LfD model with pre-training (*i.e.*, proposed) and without pre-training (*i.e.*, baseline shown in blue).

	Stiffness		Surface Friction	
	Interpolation	Extrapolation	Interpolation	Extrapolation
Demo only (Baseline)	12.5 ± 5.1	25.1 ± 22.9	9.0 ± 0.6	13.2 ± 5.0
Pre-trained sim (Proposed)	5.6 ± 2.2	6.4 ± 3.6	5.4 ± 0.4	4.4 ± 1.7

TABLE II: RMSE of force profile (N) when playing back the demonstrated and generated wiping motion for unseen objects.

decoder model for 200 epochs at a learning rate of 0.0001, using a loss function Eq. (1) with  $\beta = 0.06$ . We set the latent space dimension to be  $z \in \mathbb{R}^5$  to implicitly represent the stiffness, surface friction, and other non-intuitive physical properties of the sponges. We use the Adam optimizer [36] to update the weights of the neural network.

Then, we train the motion decoder as described in Section II-B using demonstration data produced by a human demonstrator on the real robot. The motion decoder consists of 1 fully connected layer with a dropout rate of 0.1. We train the motion decoder for 10000 epochs at a learning rate of 0.001 using a loss function Eq. (2).

#### IV. RESULTS AND DISCUSSION

##### A. Does pre-training haptic representation encoder improve the generation of desired behaviour for unseen objects?

Fig. 3 shows the Root Mean Squared Error (RMSE) between the demonstrated and generated motions for four test cases: stiffness interpolation, stiffness extrapolation, friction interpolation, and friction extrapolation. We train the model five times with randomly initialised weights. In all four test cases, we observe a statistically significant reduction in the RMSE of the generated wiping motion trajectory when we pre-train the model using 1000 unsupervised haptic data (orange), compared to when we use only demonstration data without pre-training (blue). It is important to note that the proposed approach achieves this improvement despite using the same number of demonstrations for training (*i.e.*, without increasing human effort).




Fig. 4: Plots of the wiping motion and force profile in the normal direction when playing back the wiping motion trajectory demonstrated and generated by the LfD model with pre-training (*i.e.*, proposed) and without pre-training (*i.e.*, baseline) of the haptic encoder.

In addition, the comparison of models pre-trained on different sizes of unsupervised haptic data (*i.e.*,  $M = 10, 200$  and 1000) shows that the RMSE between the demonstrated and generated motion trajectories decreases as the amount of unsupervised haptic data increases. This result highlights the importance of pre-training the haptic representation encoder on a large amount of unsupervised haptic data, enabled by the proposed decoupling of the haptic representation encoder from the motion generation decoder.

Next, we played back the demonstrated and generated wiping motion trajectories for unseen sponges on a real robot. In all four test cases, we observe a smaller RMSE of the force profiles during the wiping with pre-training, as shown in Table II, reflecting a decrease in the RMSE of the generated end-effector motions as shown in Fig. 3.

Fig. 4 shows an example plot of end-effector motion and force profile when playing back a wiping motion on a real robot using an unseen sponge with low stiffness and high surface friction in the case of surface extrapolation. In this particular example, as the surface friction is high, the human demonstrator applies less force in the normal direction to maintain a smooth wiping motion compared to the sponges used for training, which have a lower friction. The pre-trained model exhibits similar behaviour by generating a wiping motion with less force in the normal direction. In contrast, the model without pre-training generates motion that exerts more force, closer to the sponges with lower friction encountered during training. Notably, the maximum difference in force between the demonstrated and generated motion is 23N for the pre-trained model, while the model without pre-training exhibits a maximum force difference of up to 48N. The ability to adapt the motion to the sponge with higher friction, which the model without pre-training fails to adapt to, showcases the advantage of the pre-trained model in manipulating unseen objects that fall outside of the demonstrated data distribution.

##### B. How does the haptic representation encoder trained in simulation capture the properties of real objects?

Next, we employ t-SNE to visualise and analyse how the haptic encoder captures the properties of real objects in its latent space. We first present an example of t-SNE applied

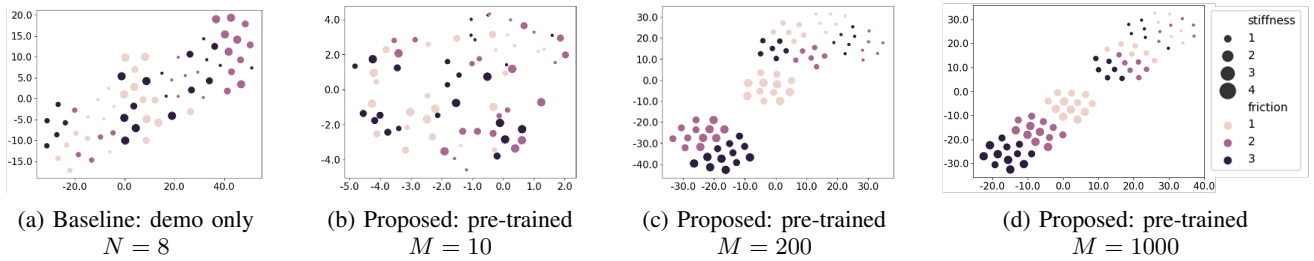


Fig. 5: t-SNE applied to the latent space of the haptic encoder when tested on the real objects. The baseline model is trained only using demonstration data. The proposed models are trained with different size of simulated unsupervised data.

to the latent space of a haptic encoder trained jointly with a motion decoder using only demonstration data (*i.e.*, baseline) to evaluate its ability to distinguish the properties of unseen real sponges in the case of stiffness extrapolation. The visualisation in Fig. 5a shows a lack of clear organisation in its latent space with respect to the stiffness and surface friction of the sponges, as 6 demonstration data (2 stiffness levels  $\times$  3 surface friction levels) are insufficient for learning to distinguish the physical properties of unseen objects.

In contrast, when we train the haptic encoder on 1000 simulated unsupervised haptic data and evaluated it on real sponges (*i.e.*, proposed), the haptic encoder is able to capture both the stiffness and the surface friction of the real objects in its latent space, as shown in Fig. 5d. The latent space exhibits a clear organisation from compliant to stiff sponges, as indicated by the size of the markers. Similarly, within sponges of the same stiffness, the latent space exhibits a clear organisation from sponges with smooth to rough surface friction, as indicated by the colours of the markers. We observe limited organisation in the latent space with respect to the surface friction of the real objects compared to the stiffness. One possible reason for this is due to the mismatch between the simulated and real-world force observations during exploration, arising from the larger sim2real gap in the friction model compared to the stiffness model.

Next, we apply t-SNE to the latent space of the haptic encoders pre-trained on unsupervised haptic data of different sizes ( $M = 10, 200$  and  $1000$ ), as shown in Fig. 5b-Fig. 5d, respectively. The results indicate that as the data size increases, the latent space becomes more organised based on the physical properties of the objects. There is a strong relation between how adequately the haptic representation encoder captures the physical properties of the objects and the RMSE of the generated motion. This relationship highlights the necessity of capturing the underlying distribution of object properties in order to generalise motions learnt from demonstrations to unseen objects.

### C. Does the haptic representation encoder capture the physical properties of objects?

Next, we train the haptic representation encoder using only each of the exploratory actions to gain further insight into the information encoded in the latent space. When using force trajectories obtained from pressing only, the latent space clearly captures the stiffness of the sponges, as indicated by the marker sizes in Fig. 6a, while the surface




Fig. 6: t-SNE applied to the latent space of the haptic encoder when performing only (a) pressing and (b) lateral motion.

friction of the sponges remains indistinct. Conversely, when performing the lateral motion only, the latent space distinctly captures the surface friction of the sponges, as indicated by the marker colours in Fig. 6b, while the separation with respect to the stiffness of the sponges is unclear. The latent space captures the properties of interest corresponding to each designed exploratory action. These results indicate that the haptic representation encoder encodes meaningful information (*i.e.*, the physical properties of the object) without explicitly specifying which physical properties to encode or providing access to the ground truth of the physical property parameters.

## V. CONCLUSIONS

Our research addresses the challenge of enabling robots to learn to manipulate objects with diverse physical properties from limited human demonstrations. We propose a semi-supervised LfD approach to pre-train the haptic representation encoder independently of the motion generation decoder, utilising unsupervised data without relying solely on demonstration data. Experimental results show that the pre-training improves the ability of the LfD model to recognise the physical properties and generate the desired wiping motion for unseen sponges, outperforming the LfD without pre-training. In addition, our analysis of the latent space of the haptic representation encoder demonstrates its capability to capture real object properties when trained on simulated unsupervised data.

In this work, we carefully chose the exploratory actions considering the physical properties of the objects relevant to the downstream task. An interesting area for future research involves generating such exploratory behaviours by observing human exploration or by optimising information gain. In addition, we would like to relax the assumption of unchanged object properties during task execution and develop methods to enable online updates of object properties and adaptation of behaviours.

## REFERENCES

- [1] F. J. Valero-Cuevas and M. Santello, "On neuromechanical approaches for the study of biological and robotic grasp and manipulation," *Journal of neuroengineering and rehabilitation*, vol. 14, pp. 1–20, 2017. DOI: 10.1186/s12984-017-0305-3.
- [2] R. S. Dahiya, G. Metta, M. Valle, and G. Sandini, "Tactile sensing—from humans to humanoids," *IEEE transactions on robotics*, vol. 26, no. 1, pp. 1–20, 2009. DOI: 10.1109/TRO.2009.2033627.
- [3] S. J. Lederman and R. L. Klatzky, "Hand movements: A window into haptic object recognition," *Cognitive Psychology*, vol. 19, no. 3, pp. 342–368, 1987. DOI: 10.1016/0010-0285(87)90008-9.
- [4] K. Kawaharazuka, N. Kanazawa, K. Okada, and M. Inaba, "Learning-based wiping behavior of low-rigidity robots considering various surface materials and task definitions," in *IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids)*, 2022. DOI: 10.1109/Humanoids53995.2022.10000172.
- [5] M. Do, J. Schill, J. Ernesti, and T. Asfour, "Learn to wipe: A case study of structural bootstrapping from sensorimotor experience," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014. DOI: 10.1109/ICRA.2014.6907103.
- [6] M. C. Gemicic and A. Saxena, "Learning haptic representation for manipulating deformable food objects," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2014. DOI: 10.1109/IROS.2014.6942626.
- [7] P. Sundaresan, S. Belkhal, and D. Sadigh, "Learning visuo-haptic skewering strategies for robot-assisted feeding," in *6th Annual Conference on Robot Learning*, 2022. [Online]. Available: <https://openreview.net/forum?id=lLq09gVoaTE>.
- [8] T. Bhattacharjee, G. Lee, H. Song, and S. S. Srinivasa, "Towards robotic feeding: Role of haptics in fork-based food manipulation," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1485–1492, 2019. DOI: 10.1109/LRA.2019.2894592.
- [9] T. Lopez-Guevara, R. Pucci, N. K. Taylor, M. U. Gutmann, S. Ramamoorthy, and K. Subr, "Stir to pour: Efficient calibration of liquid properties for pouring actions," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020. DOI: 10.1109/IROS45743.2020.9340852.
- [10] N. Saito, N. B. Dai, T. Ogata, H. Mori, and S. Sugano, "Real-time liquid pouring motion generation: End-to-end sensorimotor coordination for unknown liquid dynamics trained with deep neural networks," in *IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 2019. DOI: 10.1109/ROBIO49542.2019.8961718.
- [11] T. Davchev, K. Luck, M. Burke, F. Meier, S. Schaal, and S. Ramamoorthy, "Residual learning from demonstration: Adapting dmps for contact-rich manipulation," English, *Robotics and Automation Letters*, vol. 7, no. 2, pp. 4488–4495, 2022. DOI: 10.1109/LRA.2022.3150024.
- [12] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, "Robot programming by demonstration," in *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 1371–1394. DOI: 10.1007/978-3-540-30301-5\_60.
- [13] A. Hussein, M. M. Gaber, E. Elyan, and C. Jayne, "Imitation learning: A survey of learning methods," *ACM Comput. Surv.*, vol. 50, no. 2, Apr. 2017. DOI: 10.1145/3054912.
- [14] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard, "Recent advances in robot learning from demonstration," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, no. 1, pp. 297–330, 2020. DOI: 10.1146/annurev-control-100819-063206.
- [15] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a few examples: A survey on few-shot learning," *ACM Comput. Surv.*, vol. 53, no. 3, Jun. 2020. DOI: 10.1145/3386252.
- [16] S. James, M. Bloesch, and A. J. Davison, "Task-embedded control networks for few-shot imitation learning," in *2nd Conference on robot learning*, 2018.
- [17] N. Saito, T. Ogata, S. Funabashi, H. Mori, and S. Sugano, "How to select and use tools? : Active perception of target objects using multimodal deep learning," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2517–2524, 2021. DOI: 10.1109/LRA.2021.3062004.
- [18] A. Merckling, A. Coninx, L. Cressot, S. Doncieux, and N. Perrin, "State representation learning from demonstration," in *Machine Learning, Optimization, and Data Science*, G. Nicosia, V. Ojha, E. La Malfa, et al., Eds., Cham: Springer International Publishing, 2020, pp. 304–315.
- [19] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, Aug. 2013. DOI: 10.1109/TPAMI.2013.50.
- [20] J. Pari, N. M. Shafiullah, S. P. Arunachalam, and L. Pinto, "The surprising effectiveness of representation learning for visual imitation," *Robotics: Science and Systems XVIII*, Jun. 2022. DOI: 10.15607/rss.2022.xviii.010.
- [21] P. Florence, L. Manuelli, and R. Tedrake, "Self-supervised correspondence in visuomotor policy learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 492–499, 2020. DOI: 10.1109/LRA.2019.2956365.
- [22] I. Radosavovic, T. Xiao, S. James, P. Abbeel, J. Malik, and T. Darrell, "Real-world robot learning with masked visual pre-training," in *6th Annual Conference on Robot Learning*, 2022. [Online]. Available: <https://openreview.net/forum?id=KWCzfuqshd>.
- [23] A. Salari, A. Djavadifar, X. Liu, and H. Najjaran, "Object recognition datasets and challenges: A review," *Neurocomputing*, vol. 495, pp. 129–152, 2022. DOI: 10.1016/j.neucom.2022.01.022.
- [24] X. Han, Z. Zhang, N. Ding, et al., "Pre-trained models: Past, present and future," *AI Open*, vol. 2, pp. 225–250, 2021. DOI: 10.1016/j.aiopen.2021.08.002.
- [25] J. Bohg, K. Hausman, B. Sankaran, et al., "Interactive perception: Leveraging action in perception and perception in action," *IEEE Transactions on Robotics*, vol. 33, no. 6, pp. 1273–1291, 2017. DOI: 10.1109/TRO.2017.2721939.
- [26] L. Natale, G. Metta, and G. Sandini, "Learning haptic representation of objects," in *International Conference on Intelligent Manipulation and Grasping*, Genoa, 2004.
- [27] V. Chu, I. McMahon, L. Riano, et al., "Robotic learning of haptic adjectives through physical interaction," *Robotics and Autonomous Systems*, vol. 63, pp. 279–292, 2015, Advances in Tactile Sensing and Touch-based Human Robot Interaction. DOI: 10.1016/j.robot.2014.09.021.
- [28] B. A. Richardson, K. J. Kuchenbecker, and G. Martius, "A sequential group VAE for robot learning of haptic representations," in *The CoRL Workshop on Aligning Robot Representations with Humans*, 2022. [Online]. Available: [https://aligning-robot-human-representations.github.io/docs/camready\\_11.pdf](https://aligning-robot-human-representations.github.io/docs/camready_11.pdf).
- [29] I. Guzey, B. Evans, S. Chintala, and L. Pinto, "Dexterity from touch: Self-supervised pre-training of tactile representations with robotic play," in *7th Annual Conference on Robot Learning*, 2023. [Online]. Available: <https://openreview.net/forum?id=EXQ0eXtX30W>.
- [30] A. Drimus, G. Kootstra, A. Bilberg, and D. Kragic, "Design of a flexible tactile sensor for classification of rigid and deformable objects," *Robotics and Autonomous Systems*, vol. 62, no. 1, pp. 3–15, 2014, New Boundaries of Robotics. DOI: 10.1016/j.robot.2012.07.021.
- [31] O. Kroemer, C. H. Lampert, and J. Peters, "Learning dynamic tactile sensing with robust vision-based training," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 545–557, 2011. DOI: 10.1109/TRO.2011.2121130.
- [32] M. A. Lee, Y. Zhu, K. Srinivasan, et al., "Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019. DOI: 10.1109/ICRA.2019.8793485.
- [33] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv*, 2013. DOI: 10.48550/arXiv.1312.6114.
- [34] C. Mower, T. Stouraitis, J. Moura, et al., "Ros-pybullet interface: A framework for reliable contact simulation and human-robot interaction," in *6th Conference on Robot Learning*, 2023. [Online]. Available: <https://proceedings.mlr.press/v205/mower23a.html>.
- [35] C. E and B. Y., "Pybullet, a python module for physics simulation for games, robotics and machine learning," [Online]. Available: <https://pybullet.org/wordpress/>.
- [36] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv*, 2014. DOI: 10.48550/arXiv.1412.6980.